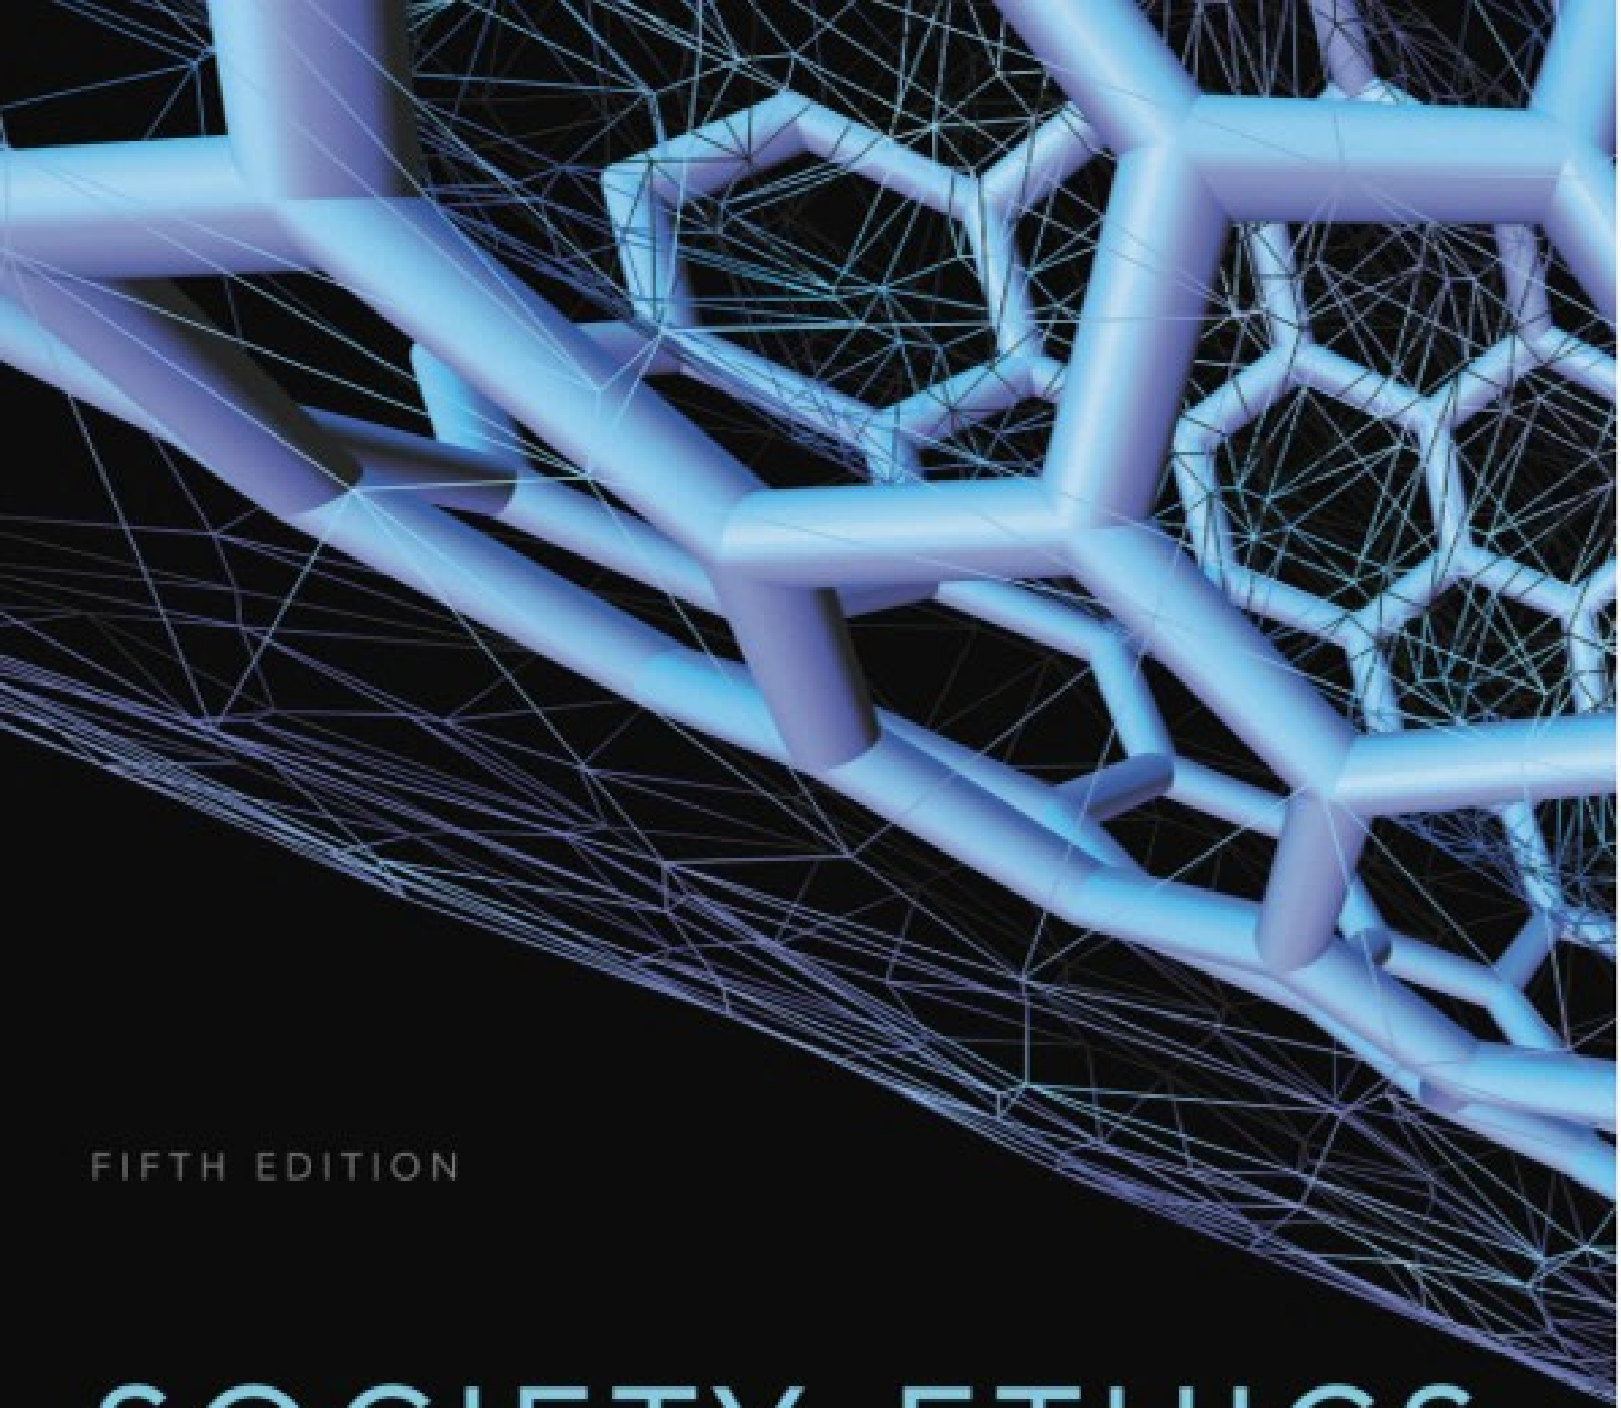

FIFTH EDITION

SOCIETY, ETHICS, & TECHNOLOGY

MORTON WINSTON • RALPH EDELBACH

CONTENTS

PREFACE xi
HOW TO USE THIS BOOK xiii

INTRODUCTION: MORTON WINSTON, Children of Invention Revisited 1

PART ONE
Perspectives on Technology 27

1.1 Historical Perspectives 28
1.1.1 JAMES BURKE: *Something for Nothing* **28**
1.1.2 RUTH SCHWARTZ COWAN: *Industrial Society and Technological Systems* 37
1.1.3 ARTHUR DONOVAN: *Containerization and Globalization: An Innovation and Its Impact* **54**
1.1.4 ROSALIND WILLIAMS: *History as Technological Change* **60**

1.2 Social/Political Perspectives 69
1.2.1 RICHARD SCLOVE: *I'd Hammer Out Freedom: Technology as Politics and Culture* **69**
1.2.2 ANDREW FEENBERG: *Democratic Rationalization* 77
1.2.3 SHEILA JASANOFF: *Technologies of Humility: Citizen Participation in Governing Science* **92**
1.2.4 JUDY WAJCMAN: *Addressing Technological Change: The Challenge to Social Theory* **103**

1.3 Ethical Perspectives 116
1.3.1 HANS JONAS: *Technology and Responsibility: Reflections on the New Task of Ethics* **116**
1.3.2 MICHAEL DAVIS: *Constructing the Professional Responsibility of Engineers* **127**

1.3.3 **DAVID STRONG:** *Technological Subversion* 145

1.3.4 **DEBORAH G. JOHNSON AND THOMAS M. POWERS:** *Ethics and Technology: A Program for Future Research* 156

PART TWO

Contemporary Technology and the Future 167

2.1 Security and Surveillance 168

2.1.1 **MAX BOOT:** *The Consequences of the Information Revolution* 168

2.1.2 **NOEL SHARKEY:** *The Automation and Proliferation of Military Drones and the Protection of Civilians* 180

2.1.3 **JAY STANLEY AND BARRY STEINHARDT:** *Bigger Monster, Weaker Chains* 189

2.1.4 **JAMES STACEY TAYLOR:** *In Praise of Big Brother* 203

2.2 Artificial Intelligence and Robotics 216

2.2.1 **ABIGAIL TUCKER:** *Birth of a Robot* 216

2.2.2 **FIORELLA OPERTO:** *Ethics in Advanced Robotics* 224

2.2.3 **STEPHEN BAKER:** *Final Jeopardy* 234

2.3 Nanotechnology 242

2.3.1 **RAY KURZWEIL:** *Neuroscience, Nanotechnology, and Ethics: Promise and Peril* 242

2.3.2 **LYNN L. BERGESON:** *Emerging Nanomaterial Governance Systems: The State of Play* 251

2.3.3 **ARMIN GRUNWALD:** *Nanotechnology—A New Field of Ethical Inquiry?* 264

2.4 Internet and Social Media 276

2.4.1 **NICHOLAS CARR:** *Is Google Making Us Stupid?* 276

2.4.2 **REBECCA MACKINNON:** *Corporate Censorship* 283

2.4.3 **ANDREW J. FLANIGAN, CRAIG FLANIGAN, AND JON FLANIGAN:** *Technical code and the Social construction of the Internet* 292

2.5 Biotechnology 306

2.5.1 **FRANCIS FUKUYAMA:** *Our Posthuman Future* 306

2.5.2 **LEON R. KASS:** *Preventing a Brave New World* 331

2.5.3 **MICHAEL J. SANDEL:** *The Case against Perfection* 343

2.6 Energy and the Environment 355

2.6.1 ROBERT W. KATES: *The Nexus and the Neem Tree* **355**

2.6.2 DONALD AITKEN: *Global Warming* **370**

2.6.3 DAVID FRIDLEY: *Nine Challenges of Alternative Energy* **386**

2.6.4 LESTER BROWN: *Stabilizing Climate: An Energy Efficiency Revolution* **397**

APPENDIX A: TECHNOLOGY AND ETHICS IN THE NEWS 415

APPENDIX B: CODES OF ETHICS 435

BIBLIOGRAPHY 455

PREFACE

This book aims to provide students with an understanding of the profound role that science and technology play in shaping modern society, as well as the social and environmental issues that arise from technological change. As in earlier editions, our primary goal has been to empower readers to become better-informed "technological citizens," and to participate in the ongoing conversations concerning the impact of new technologies on the ways we live and work in the twenty-first century. In the few years since the publication of the fourth edition, important conversations have arisen concerning topics such as drone warfare, nanoscale materials, social media, gene patenting, artificial intelligence, robotics, Internet freedom, and alternative energy technologies.

In preparing this fifth edition of *Society, Ethics, and Technology,* we have endeavored to preserve the selections and features of earlier editions of the text that instructors and students told us were most valuable, while also expanding our coverage of evolving debates concerning the benefits and risks to society of contemporary technologies.

The basic organization of the book remains the same: Sections 1.1 through 1.3 contain readings addressing the history, sociology, and philosophy of technology, with new readings by authors such as James Burke, Arthur Donovan, Andrew Feenberg, and Michael Davis. The second part of the book, Sections 2.1 through 2.6, has been reorganized and expanded in order to address social and ethical issues arising from rapid technological changes in fields such as security and surveillance, nanotechnology, Internet and social media, artificial intelligence and robotics, genetic engineering, and alternative energy. More than half of the readings in the second part are new for the fifth edition and include the work of authors such as Donald Aitken, Stephen Baker, Lester Brown, Nicholas Carr, Francis Fukuyama, Rebecca MacKinnon, and Abigail Tucker.

As in earlier editions, each selected reading has a brief introduction, a list of discussion questions, and a set of keywords. We have updated and revised these pedagogical aids for this edition. We have also included two appendices in this edition: Appendix A contains a selection of stories in the news that highlight some of the ethical and social issues of contemporary technology. These short readings are also

accompanied by questions for discussion and keywords that can be used to search for current news stories on the topic or issue. Suggestions for further reading on each topic are included to assist students in extended research. Appendix B contains a selection of ethical codes that can be used in discussions about the professional responsibilities of engineers as well as the more general responsibilities of technological citizens in a global society.

This book also has a website that provides several resources to help students and instructors alike. For students, there are case studies and a list of reliable websites on a wide variety of relevant topics. For instructors, there is a protected test bank as well as handouts used by the authors in their own classes. The website may be found by looking up the book at www.cengagebrain.com.

We wish to thank Joann Kozyrev for her support and encouragement in developing the plan for the fifth edition. We also wish to thank the following reviewers: Maureen Allen, Elon University; Suzie Cashwell, Herzing University Online; Marcel Daguerre, CSU Chico; Zarreen Farooqi, University of Akron; John Fielding, Roger Williams University; Jason Ford, University of Minnesota; Yasha Karant, CSU San Bernardino; Jil MacMenamin, Community College of Vermont–Brattleboro; Alison Miyauchi, Alberta College of Art and Design; Thomas Oberdan, Clemson University; Philip Reed, Canisius College; Philip Rose, University of Windsor; John Trono, Saint Michael's College–Vermont; Heather Wilson, Northern Essex Community College.

We are grateful to the many instructors and students who provided feedback on the previous editions, as well as useful suggestions for new topics and readings to include in this new edition. We are also grateful for the assistance and professionalism of our editorial team at Cengage, particularly the indefatigable Ian Lague, and for the expert permission requesting and the pre-production work by PreMediaGlobal in Chennai, India. Finally, we wish to thank our wives for their companionship and support while we worked to produce this new edition of our book.

MW & RE
December, 2012

THIS BOOK IS DESIGNED for use in standard fifteen-week undergraduate Science, Technology, and Society (STS) courses. Courses of that type are intended to provide an interdisciplinary bridge between the humanities, particularly ethics, the social sciences, and the natural sciences and engineering by developing a framework for analyzing the social, environmental, and ethical implications of contemporary science and technology. A principal goal of these courses is to empower students to think critically about contemporary technological issues, such as privacy on the Internet, genetic engineering, and global warming, and to learn to accept the social responsibilities of educated citizens in a global technological society.

The key bridge idea that links science and technology with ethics and social science is the dictum that "ought implies can, but can does not imply ought." The first half of this statement asserts that one cannot be obliged to do something that one cannot do; the second half says that just because one can do something does not mean that one ought to do it. Science and technology, by enlarging the domain of what we humans know how to do, are continually enlarging the class of actions that we can perform. As a consequence of this expansion of our power to do more things, we are obliged to ask ourselves whether we ought to be doing all the things that science and technology now allow us to do. "Ought questions" are ethical questions, and to answer these sorts of questions, one needs to understand how ethical values and traditional moral standards can be applied to evaluating past and present scientific and technological innovations.

The introductory essay, "Children of Invention Revisited," provides students with a conceptual framework for thinking about the relationship between scientific and technological innovation and the social and ethical issues that such innovations can raise when new technologies are deployed and diffused in society. All too often, new technologies are developed and deployed without proper reflection on their likely consequences for society. One reason for this is ignorance about the nature of technology and the processes by which new technologies are developed and deployed. Addressing the need for students to be scientifically and technologically literate has

long been a primary goal of the STS movement, and it is also a primary goal of this book.

However, another reason for the failure to evaluate technology ethically is that traditional ethical theories were primarily intended to evaluate persons and their actions rather than socio-technical practices. To correct this, students need a little understanding of ethics and to be familiar with the essential features of the major traditional ethical theories, such as ethical egoism, utilitarianism, natural law, and Kantian deontological ethics. Explaining how these theories can be adapted to the evaluation of technology is also necessary so that methods are provided for applying familiar ethical principles to particular controversial issues surrounding contemporary science and technology. Finding books that serve these multiple purposes is unusual.

The introductory essay attempts to provide an ethical framework through which students can approach the evaluation of technology. We advise that the introductory essay be assigned early in the course because it provides students who are new to STS with a comprehensive overview of the subject matter as well as a concise introduction to the ethical and social aspects of technology. The selections in Part 1 convey general perspectives on the historical, social, and philosophical aspects of technology, and those in Part 2 deal with several issues in specific fields of contemporary technology.

In selecting readings for this text, we have tried to provide current selections on a number of standard STS topics, such as computers, automation, robotics, the Internet, genetic engineering, artificial reproduction, population, energy, and environmental policy issues.

Each selection is preceded by a brief introduction describing the topic and major themes of the reading, mentioning something about its author, and providing several focus questions for students to keep in mind while reading the essay. Usually, at least one of these focus questions asks the reader to compare the view in the present reading with that found in another selection that presents a competing or complementary perspective. This allows instructors to design their syllabus using such topical links instead of the standard organization.

However, one of the problems with edited readers is that no matter how careful the editors are in selecting the texts to be included, it is inevitable that some topics or issues will be omitted or given only cursory treatment and some readings will soon become less relevant. This problem is especially difficult with respect to the field of science, technology, and society studies because, given the dynamic nature of contemporary science and technology, new discoveries and developments are occurring constantly and their impacts on society need to be addressed and examined. Thus, even if the selections of an edited anthology are current at the time of publication, they soon lose their currency and must be updated with new materials.

The companion website associated with this text, which can be easily located by searching for the title of this book at CengageBrain.com, can be of value to both students using this book in a class or instructors teaching this type of course. Because information in this area is constantly being updated and expanding, it is helpful to have

one source that will bring together many different types of related information. The website is designed to do that plus provide an opportunity for both students and faculty to share ideas, information, and additional teaching resources. If you have materials that you would like to share with others to be included on the website please contact us.

Morton Winston
Ralph Edelbach

INTRODUCTION

Children of Invention Revisited

Morton Winston

IF NECESSITY IS THE MOTHER of invention, who is the father, and who, or what, are invention's children? Necessity, of course, is a matter of degree: We actually only *need* air, water, and food to survive. Shelter, clothing, and a few material possessions are also nice, as are companionship, affection, security, and several other psychological goods that we crave as social animals. But we humans learned how to satisfy these basic biological needs millions of years ago. Why then did we embark on the long journey that transformed us from cavemen into cosmonauts? What was it that made possible the ascent from the Stone Age to our present global technological civilization?

Clearly, the leading answer to these questions is superior intelligence. But in what respects is human intelligence superior to that found in other species? Is it our capacity to learn from observation and experience and to transmit what we learn to others? Is it our ability to create and use language? Or might it be these two general cognitive capacities for culture and language, together with our unique ability to discover new solutions to old problems, better ways of making and doing things? In short, is it our unique capacity as a species to form science and technology? Coupled with our needs and desires, which provide the motives that propel us to discover and invent, our scientific and technological creativity has guided the development of civilization through the development of theories, tools, inventions, and technologies that have transformed the ways that we live and work.[1]

For most of us, a world without technology is inconceivable. The inventions that it has given us are all around us. In fact, most of us spend most of our lives in completely artificial environments, wrapped in a technological cocoon that provides us with much more than merely food and protection from the elements. We are so wrapped up in our technological culture, in fact, that it takes an effort to distance ourselves from it in order to understand how technology has transformed human existence from its natural state. Such a historical perspective also helps us see how contemporary technologies, such as genetic engineering and the Internet, are now changing us in even more dramatic ways, creating new opportunities for humans to flourish, new ways of life, and also, in some cases, new social and ethical problems. These social and ethical issues arising from technological innovation are the "children

For our children and grandchildren

FIFTH EDITION

SOCIETY, ETHICS, AND TECHNOLOGY

MORTON E. WINSTON
The College of New Jersey

RALPH D. EDELBACH
The College of New Jersey

WADSWORTH
CENGAGE Learning

Australia • Brazil • Japan • Korea • Mexico • Singapore • Spain • United Kingdom • United States

appeared around 1800 in Paris. Early "velos" were powered by pushing on the ground with the feet and were not steerable until 1818. Velos were powered using treadles and rods by 1840 and by foot pedals in 1861. Finally, in approximately 1869, a chain was used to transmit power, creating a machine that is recognizable as a modern bicycle.

In other situations, an invention's date was established by determining when the device or process was first depicted in illustrations or mentioned in writings. In yet other cases, the date a patent was issued for the process or device is indicated. Obviously, the dates of recent discoveries are accurate as a result of better record keeping and documentation.

In some cases, two dates are shown matching the two places that those events/discoveries (or variations of the technology) took place. This may be useful in understanding the time that elapsed for a given technology in moving from one locale to another (technology diffusion) or in evolving separately in two different societies.

Event	Date
Chronometer	1735
Circumnavigation of globe	1589
Cities	8000 BCE
Clock (mechanical)	1300
Clock (pendulum)	1656
Clock (water)	1400 BCE
Cloning (fish)	1981
Coca-Cola	1866
Codes/ciphers	2000 BCE
Coffee	1400
Coins	700 BCE
Compact discs	1982
Compass	300 BCE
Computer (mechanical)	1832
Computer (electromechanical)	1942
Computer (electronic)	1946
Computer (digital/binary)	1938
Computer (personal)	1975
Computer (programmable)	1941
Concrete	200 BCE
Contact lens	1887
Container shipping (ships/trucks)	1954
Contraception	2000 BCE
Contraceptive pill	1956
Contracts (written)	2700 BCE
Copper	6400 BCE
Copying (Xerography)	1937
Crank handle	100 BCE
Crayons (Crayola)	1902
Crop rotation (theory)	1804
Crossbow	200 BCE
Crystal radio	1901
Diaper (disposable)	1951
Dinosaurs (era)	245 to 65 MYA
DNA discovered	1953
DNA evolves	1.8 BYA
Drill (bow)	20,000 BCE
Drill (dental)	1790
DVD	1995
Dynamite	1867
Earth Day	1970
Electrocardiogram (EKG)	1924
Electroencephalogram (EEG)	1929
Elevator (electric)	1887
E-mail	1987
Embryo transplant (cattle/human)	1950/1983
Erector set	1901
Escalator	1911
Fastener (Velcro)	1948
Fastener (zipper)	1923
Fax machine	1902
Fetal surgery	1984
Fiber optics	1955
Fire (human made)	1.5 MYA
First mammals	65 MYA
Fork	600
Formation of Earth	5 BYA

(continued at back of book)

of invention" that this book is about. To understand these issues, however, it is first necessary to get a clear view of their source—*technology*.

THE SCOPE OF TECHNOLOGY

The word *technology* is itself of fairly recent coinage; Johann Beckman of Gottingen first used it in 1789. Its root, *techne*, is the ancient Greek word for "art," "craft," or "skill," which itself is derived from an earlier Indo-European root, *teks*, which means "to weave" or "to fabricate" (*teks* is also the root of the word *textile*). Recent archeological evidence suggests that the weaving of cloth predates the birth of agriculture and the dawn of civilization, going back to about 35,000 BCE, making it one of the first technologies. As the etymology suggests, a *techne* is a method, craft, or skill used in making things, not the things themselves, which are called *artifacts*. For instance, a woven object made from animal hairs that have been twisted together into long strands, dyed with vegetable colors, and interlaced by a weaver is an artifact. Let's say that this object functions primarily as a blanket; a person wraps her- or himself in it to stay warm. A typical use, or function, of an artifact is called its purpose or end, and the knowledge of how to gather the fibers, twist them, dye them, and weave them are the individual *techniques* that comprise this particular technology. Thus, the core meaning of the word *technology* refers to the ensembles of techniques by which humans make artifacts that serve certain useful ends. However, this original meaning is too restrictive for the contemporary context in which we think about the relationship between technology and modern society.

In recent years there has been an unfortunate tendency to narrow the definition of technology to contemporary information-communications technologies (ICTs) such as personal computers, the Internet, and the digital gadgets advertised in *Wired*. This way of thinking about technology is clearly too restrictive; it ignores other areas of contemporary technological innovation, such as biotechnology and nanotechnology, as well as the technologies of earlier periods, such as the automobile, the steam engine, or the water wheel. When one ordinarily thinks of technology, what most likely comes to mind are technological *artifacts*—the objects, machines, structures, and devices that are the useful end products of technological design. Then, perhaps, one thinks of the less familiar but potentially more impressive machines and industrial processes tucked away in the factories that manufacture the various gadgets and widgets that we use. Finally, one might visualize the scientists, engineers, and technicians in white laboratory coats, hard at work in the laboratories of the Research and Development Division, designing the next generation of technological devices and processes.

Although it is true that each context through which artifacts come into being—design, manufacturing, and end use—is a technological context, it is still too narrow a view to identify technology with only the material culture of designed or manufactured physical objects. We take an even broader view: Technology consists of not only useful artifacts and the tools and processes needed to produce them but also the entire social organization of people and materials that permits the acquisition of the knowledge and skills needed to design, manufacture, distribute, use, repair, and eventually

Technology includes not only the obvious candidates—the mechanical, structural, and electronic know-how that directs the purposeful organization of materials—but also the less obvious *invisible technologies* that control the purposeful organization of people and their labor. Mechanical clocks enabled people to coordinate their activities and thus made possible a more productive use of human labor. But clocks and calendars are useful as ways of measuring units of time—the minutes, days, weeks, months, and the like—that comprise the invisible technology of time. The monetary system, the banks, and the stock and commodity markets are technologies for the distribution of economic value that was once associated with gold coins, then with pieces of paper, currency notes, or stock certificates, and is nowadays represented by encrypted bits of digital data. The ideologies of free-market capitalism and centralized planned economies are competing economic theories about how best to organize social production. Even governmental systems, ranging from varieties of representative democracy to theocracy and dictatorship, are competing political technologies for managing concerted societal action and resolving political conflicts. People ask, "Is there a better way to run the government?" no less frequently than "Is there a better way to design a mousetrap?" Both questions are requests to find a better technology—that is, to acquire knowledge or skill that enables one to solve a practical problem.

Contemporary writers often speak of technology as consisting of systems, for instance, the telegraph and telephone, the railroad, the petroleum, and the electrical systems that came about in the later half of the nineteenth century. Large-scale technological systems are linked with one another, often in relationships of mutual interdependence; for instance, telegraph wires were strung along railroad rights of way, and railroads came to depend on the telegraph for scheduling and signaling. Similarly, contemporary ICT systems, such as the Internet, depend on a great many other technological systems for their creation and use but then are used by them, creating a matrix of complex interdependencies. One might think of the entire *technosphere—that* is, the sum total of all human-created artifacts together with the enabling knowledge that created it and sustains it—as constituting one giant technological system. However, this definition of the scope of technology is too broad to be of much practical use. Instead, we will think of technologies as consisting of several distinguishable but interacting aspects: (1) skills, techniques, human activity-forms, or sociotechnical practices; (2) resources, tools, and materials; (3) technological products, or artifacts; (4) ends, intentions, or functions; (5) background knowledge; and (6) the social contexts in which the technology is designed, developed, used, and disposed of. These six aspects are present in every technology.

The first aspect of technology is the *human activity-form*—that is, the particular skills, techniques, methods, practices, or ways of doing things. We know that animals other than humans can make and use tools; for instance, chimpanzees strip branches off tree limbs to make sticks that can be used for gathering insects. For the purposes of our characterization of the technological system, we restrict activity-forms or techniques to those employed by human beings. Some human activity-forms employ natural objects rather than tools to achieve ends; for instance, if one throws a spear in order to try to kill an animal for food, one is employing a particular technique. But throwing spears is a primitive and not very useful hunting technique; our technologies for providing our food have improved considerably. Today, there are complex ensembles of

leveraged hostile takeover, from cooking lasagna to programming a computer to sort sales data. Such complex techniques represent what is called *procedural knowledge*, or more commonly "know-how," and is contrasted with *propositional knowledge*, or "know-that." Both of these types of knowledge are necessary aspects of technological systems, but techniques are its essence. Procedural knowledge forms the basis of technology because it provides the patterns for the sociotechnical practices or human activity-forms that we use to create artifacts of all kinds and to build and maintain our complex technological systems.

One of the main consequences of technology is to increase our capacity to do things. Technologies, techniques, and tools extend, enhance, and sometimes even replace our natural powers such as sight, hearing, muscle, and even memory and thought. By using tools, we can accomplish things that we could not otherwise achieve and to do things that we could not otherwise do, thus increasing our repertoire of human activity-forms. Tools are artifacts at our disposal that can be used to make other artifacts, but tools, even the dawn stones used by our distant ancestors, are themselves artifacts that have been transformed from their natural states in some way by means of human action.

Earth itself is of course not an artifact but has for many centuries been viewed as a *resource well* into which we can dip at will in order to satisfy our needs and desires. Technology requires resources of various kinds as inputs to technological processes, and by employing specific techniques or human activity-forms, we act on and transform these resources from their original or natural states. Once a *built environment has* been created, however, everything in it can serve as a resource to further technological development. The term *infrastructure* describes elements of the built environment that are available to be used to create or apply new technologies. We live on an increasingly anthropogenic planet, one in which the evidence of the built environment can be seen from outer space in the form of clusters of light emanating from our major cities. In fact, if we include the unintended effects on Earth's atmosphere and climate caused by anthropogenic global warming, there are very few things on Earth that are unaffected by human activity.

By acting on either natural or artificial resources, through techniques, we alter them in various ways and thus create *artifacts*, which form the third aspect of technologies. A clay pot is an example of a material artifact, which, although transformed by human activity, is not all that far removed from its natural state. A plastic cup, a contact lens, and a computer chip, on the other hand, are examples of artifacts that are far removed from the original states of the natural resources needed to create them. Artifacts can serve as resources in other technological processes. This is one of the important interaction effects within the technological system: Each new technology increases the stock of available tools and resources that can be employed by other technologies to produce new artifacts, forming the *artifactual platform*.

The fourth aspect concerns the ends or functions of an artifact or technique. Most artifacts have typical or intended uses, but artifacts can in fact be embedded in multiple contexts of use or can serve multiple ends, a property called *polypotency*. However, most artifacts have an intended use, *or focal function*; a toaster, for instance, is designed to lightly burn slices of bread, but it is also polypotent and can be used as hand warmer or as a murder weapon. There is a double ambiguity in the relations

can be used to accomplish the same ends. For instance, I could have written this sentence with a quill pen, a pencil, a ballpoint pen, a typewriter, or a personal computer (PC) running text-editing software (although I used the last). And I could have used my PC to play an adventure game or calculate my income tax instead of writing this sentence. Because artifacts are designed and created to serve certain functions, it is possible to talk about the ends of these objects—that is, their intended purposes or focal functions even though the objects themselves may often also be used in ways that were not intended. The term *valence* is sometimes used to refer to the typical or conventional uses of artifacts, which may or may not match their intended purposes.[3]

The fifth aspect of technological systems is *knowledge-that*, or factual knowledge about what the universe consists of and how it operates. To employ our technologies, we need background knowledge of various kinds: what resources to use and where to find them, what techniques to employ to fabricate various artifacts, the ends and purposes that are typically served by various techniques and objects, and how all these elements fit together in a systematic way. Both knowledge-how and knowledge-that have always been an important aspect of technologies. However, since the scientific revolution of the seventeenth century, scientific knowledge—that is, both factual and theoretical knowledge about the universe and the way it works—has come to play an increasingly important role in technological development. Modern technology is inconceivable without the Scientific Revolution, and basic scientific research is continuing to open up new technological possibilities.

The sixth aspect of technology is the *social context* or organization in which technologies are developed, distributed, and employed. A division of labor in which different individuals perform different tasks or occupy different roles to accomplish common or coordinated ends characterizes technological societies. The schemes that we use for organizing human labor represent a kind of technology that can be applied to the most important resource of all—ourselves. Complex schemes for organizing human activities that have become more or less institutionalized can be called *social artifacts*. Examples of social artifacts include the stock market, battalions or divisions in an army, baseball teams, hospitals, schools, and corporations. In each case, human resources are organized in a particular way according to a plan or technique involving a division of labor in which different persons occupy different roles, and their labor is coordinated to accomplish specific sorts of goals. It is important to understand that technology encompasses not only material artifacts but also social and organizational forms and even the cognitive techniques that produce the material and social infrastructure of human civilization. These *invisible technologies* frequently consist of formal, mathematical, or analytical techniques—for instance, the scientific method, statistical analysis, or procedures for creating a balance sheet—and many other specific, high-order thinking skills, which are the content of higher education. Becoming a scientifically or technologically educated person consists mainly in the acquisition of a fairly extensive repertoire of such cognitive techniques.

The social and psychological aspects of technological systems are the least obvious but also the most important. Technology is a human social construction. This is true in an obvious and straightforward sense when we speak of large technological structures—such as bridges, buildings, or dams, which obviously came into existence only by the coordination of the activities of numerous individuals—but it is equally true in the case of the lonely amateur inventor toiling in the attic. Inventions today are

techniques employed and the knowledge by which they are put to use by the inventor are themselves the products of prior social processes. Even the inventor's own knowledge and abilities have been shaped by his or her education and by the repertoire of cognitive techniques that he or she has acquired through education. So, there is really very little, only the raw materials and the laws of nature, that has not in some way resulted from a process of social production. Even when an inventor succeeds in inventing something new, it is still unlikely to be brought into production and placed on the market unless it has some social value or is of use to other people. So, all technologies must be seen as embedded in social contexts of development, deployment, and use.

To summarize this discussion, we can define technological systems as the *complex of techniques, knowledge, and resources that are employed by human beings in the creation of material and social artifacts that typically serve certain functions perceived as useful or desirable in relation to human interests in various social contexts.*

TECHNOLOGICAL REVOLUTIONS

The use of technologies to satisfy our needs is a fundamental feature of human nature. All human societies we know of, both those presently existing and those that existed hundreds of thousands of years before the dawn of civilization, were technological to some degree. For almost all of our species' evolution, we lived in small, nomadic bands whose main means of livelihood were hunting, gathering, and scavenging. But we were also toolmakers and tool-users during this long period of human evolution, and tools were the principal means by which we satisfied our physiological needs for food, warmth, and shelter. Our hominid ancestors first began chipping stones to make simple hand tools about 2.5 million years ago. Fire was used as early as 1.5 million years ago. *If Homo sapiens* (literally, "man the wise") is now the dominant species on the planet, it is in large part because he is also *Homo faber* ("man the maker").

Early human societies were organized as hunter-gatherer groups, gathering edible plants in season and supplementing their diet with the meat or marrow of hunted animals. Quite likely, these bands of hunter-gatherers were nomadic, following animal migrations and seasonal food-plant distributions. As with present-day hunter-gatherers, ancient nomadic societies were severely limited to only those objects that they could take with them; thus, they tended to develop simple portable technologies for hunting, gathering, cooking, transportation, and defense. Perhaps surprisingly, life does not seem to have been especially hard for hunter-gatherers. The secrets of their success seem to have been populations that did not exceed the food supply, simple and limited material needs, and the ability to move to another area when the local food supply ran out. Nomadic hunter-gatherer societies have persisted into the twentieth century in such diverse environments as the African desert, the tropical rain forest, and the Arctic tundra. Remoteness might be the key to avoiding conversion to more technologically intensive ways of life. For the rest of us, our lives now deeply depend on far-flung and complex technological systems.

About 10,000 years ago, the first great technological revolution occurred in several fertile river valleys of Asia Minor and North Africa. During the *agricultural revolution*, humans learned how to domesticate animals and to plant, grow, and harvest crops to sustain their existence. This enabled humans to give up the nomadic

astronomy, class structures, patriarchy, and other social institutions that have since come to characterize the human condition. With the adoption of settled agriculture in the fertile river valleys, the history of humankind begins. Permanent houses could be built, tools and objects could be accumulated from year to year, and so humanity began the long climb toward the collections of miscellany and junk that now clutter people's closets, attics, and garages.

Settled agriculture had many advantages and a few disadvantages. The quantity of food that could be produced per acre was much higher, so population densities could also be much greater. With permanent dwellings, creature comforts could be made that did not have to be portable. With larger numbers of people living together, specialization of activities could take place, and specialists were more likely to find better ways to do things. Larger concentrations of people could better share and perpetuate knowledge and band together to cooperate on projects that smaller groups could not attempt. Thus, we see that even at this early stage of technological development, the organization of people, information, and accumulated resources were essential aspects of emerging technological societies.

In regions with insufficient rainfall to sustain many crops, it was necessary to design, construct, and maintain either irrigation canals or aqueducts. There is evidence of canal irrigation in both Mesopotamia and Egypt as early as the sixth millennium BCE, and in areas where the topography posed challenges various devices were developed to raise water above its natural level. Some of these devices, such as the noria used with flowing water, were sophisticated; others, such as the chain-pump used with still water, were simple, being powered either by animals or humans. Devices of the latter type are still being used today in some parts of the world. Even with the Nile River's normally adequate supply of water for irrigation in Egypt, it was usually necessary to employ technology to direct and control its distribution, making agriculture a more complex undertaking than originally might be thought.

The disadvantages of settled agriculture sprang from the fact that society had "put all its eggs in one basket" and had committed itself to living in one place. A settled society is prey to flood, drought, and insects. Persistent weeds must be removed from fields before they displace crops. Houses and farm implements must be maintained. Crop seeds must be gathered and sown. The final product, food, must be harvested, stored, and distributed. In short, the settled farmer has more but must work harder to maintain his or her improved standard of living. Irrigated agriculture is even more technologically intensive and requires more complex social organization. Large irrigation projects demand larger groups to support them and must be maintained throughout the year, not just during the growing season. Irrigated farms produce more food per acre, more reliably than dry farms that rely on uncertain rainfall, but they also require more work per person fed. At the extreme are rice paddies in the river deltas of southeast China where three crops are grown each year. They are the most productive farmlands but also the most labor intensive. Today, most agricultural production in industrialized countries occurs on large farms where energy-intensive farm machinery substitutes for human labor and chemical fertilizers maintain soil fertility.

The second great technological revolution took place many centuries later, during the eighteenth century in Europe about 300 years ago. The *Industrial Revolution* replaced the muscle power of animals with coal-fired steam energy and then later

England, but before long they were improved and used to power looms and other machines in factories. The machine age caused profound changes in economic and social relations. The number of people needed to produce food declined as the number of people engaged in factory work increased. People migrated from rural areas to cities in search of higher-paying factory jobs, and new inventions such as the cotton gin, the locomotive, and the telegraph laid the groundwork for the emergence of the complex technological society that we live in today.

The methods that a society uses to produce goods have a profound effect on what life is like in that society, for both producers and consumers of goods. Prior to the industrial age, production was organized by crafts. Individual artisans both designed and produced each individual product, usually guided by traditional techniques that were occasionally modified by creative innovations. The relative value of the product was largely determined by the artisan's skill. As a result, artisans were relatively autonomous, and production units often consisted of a single artisan and several apprentices in cottage industries.

When the invention of the steam engine made power available on a scale never previously possible, it became feasible to concentrate larger numbers of workers in one place, and to have each worker perform only a small part of the production process. This resulted in a much more specialized division of labor, and the factory system was born. The factory system required far greater concentrations of power, labor, and raw materials than either agriculture or cottage industries. It also required the development of infrastructure for transportation of raw materials to the factory site and finished products from the site. Railroads and canals were thus as essential a part of the Industrial Revolution as the factories themselves. The industrial system also required a large labor force near the factory, so society's living patterns were reorganized to include factory towns where workers lived and the means to supply them with food and other necessities. Factories were often located near sources of power or raw materials, resulting in net population shifts away from agricultural lands.

In the early twentieth century, technological experts working under the banner of "scientific management," developed by F. W. Taylor in 1911, studied the production process and learned what each worker knew about making the product. They then ordained the perfect way to produce a given product using standardized parts, the division of labor, and mass-production techniques, what each worker would do, and at what pace he or she would do it. Each worker needed fewer skills and could be paid less per item. Cheaper workers making larger numbers of products using specialized machinery resulted in less expensive goods. Lower prices resulted in increased standard of living for consumers. Factory work may have become onerous, but a salary could buy more than it could previously. In recent decades, much of the world's production has moved to low-wage countries such as China and India where workers are paid far less for their work than workers in developed countries would be paid for comparable labor, in part because workers in these countries usually do not have the right to form unions and bargain collectively with their employers. But jobs with low wages and limited rights, many claim, are often better than no jobs at all or trying to scrape out a subsistence living on small farms.

In the search for increased productivity, working conditions in early factories were often harsh and dangerous. In response to the many abuses that existed, employees

such organizing drives secured improved working conditions and raised the standard of living of millions of workers and their families. A similar process of humanizing conditions of factory workers is now taking place in the developing countries where most current production is located. Despite the dominance of the factory system, crafts did not vanish entirely. They survived in niches where no one could think of an economical way of applying mass-production techniques or as a way to produce distinctive, high-quality goods. In some cases, they survived because traditional cultural values prevailed over the lure of newer technologies. In many countries sewing machines continued to be used in homes to make clothes for the family, and in India, Mohandas Gandhi revived the spinning wheel as an alternative to mass-produced thread. But the dominant trend throughout the latter half of twentieth century was toward mass-produced, globally distributed consumer goods produced by workers in low-wage countries.

The technologies of power production were driving forces of the industrial system, and each new source of power required industrialized society to provide an accompanying infrastructure to make the system work. Water power, an ancient technology, was limited in availability and location prior to the building of aqueducts, and required relatively little additional infrastructure beyond that already available in an agricultural and craft society. Coal could be more widely distributed, but coal-powered factories were large because efficient steam engines were large. Railroads and canals began to crisscross the countryside from mine to factory to market. Monetary supply and financial services had to expand to serve a system with increasing separation between producer and consumer. Electricity was later adopted as a more flexible source of power, capable of efficiently driving both large and small machines. Electricity permitted greater decentralization of industry supported by a network of power grids that eventually reached nearly every house and factory in the country. Oil and gasoline revolutionized transportation and distribution of goods. Internal combustion engines fueled by gasoline and diesel oil made it possible to have smaller vehicles, and smaller vehicles continued the trend toward decentralization. However, gas- and oil-powered vehicles required more and better roads. The U.S. interstate highway system, built in the 1960s–1980s (and similar systems in other industrialized countries) are society's most recent contributions to an industrial technology system based on oil, a system that may now be reaching its final phrase, perhaps to be replaced during the twenty-first century by a "greener" energy system that relies more on renewable forms of energy.

Many people believe that since the mid-1970s we have been going through a third great technological transformation—from the machine age to the *information age* (also called the "third wave" and the "knowledge revolution"). Computers, communications satellites, fiber-optic cable, and other developments—which make possible global, high-capacity, high-speed communications technologies such as the Internet—are already profoundly changing the way that we live, work, and play.

Revolutionary developments in computing and communications technology have transformed the workplace, faster than some would like but slower than its visionaries had hoped for. The earliest successes of computers in industry were in payroll, inventory, and similar routine and repetitive kinds of record keeping. The automated processes were well understood, straightforward, and implemented exactly as they had

The next stage gave decision makers more and better information to enhance efficiency, competitiveness, and other factors reflected in the bottom line. Computers made it possible to gather and organize data on an unprecedented scale. Also successful were the attempts to use computers to improve scheduling and reduce inventory in the production process. Goods stored in inventory cost money to store and contribute nothing to profit until they are used or sold. Predicting exactly how much of which raw materials and parts are needed at which steps of the manufacturing process and scheduling their arrival in the factory at precisely the right place at precisely the right time was the just-in-time manufacturing technique developed in Japan that led to real gains in productivity that drove the global economy in the 1990s.

As computers and computer programmers got better, computers became capable of doing jobs that were formerly thought to require human intelligence. Typically, computers proved capable of doing far more than most people would have predicted in advance and far less than their most vocal proponents claimed was possible. Although the conceptually most impressive achievements were in areas like expert systems for medical diagnosis, the biggest successes of computer technology were in the simpler applications now so common that we take them for granted: automatic pilot, antilocking brakes, electronic fuel injection, and most important, in more flexible, general-purpose tools and machines for making other products.

With flexible, modifiable, reprogrammable tools, it was no longer necessary to have long production runs to amortize the setup time of the machinery. Computer-controlled machinery could switch quickly from one task to another, and customized production runs became in some cases economically viable. Supply could now more accurately follow demand, and both idle machinery and unproductive inventory were virtually eliminated in those industries adopting the new technology.

The synergies created by computers, user-friendly software applications, satellite-mediated communications, the Internet, containerization, and rapid and relatively inexpensive air freight made possible the kind of geographically distributed production systems that are characteristic of the contemporary era of *globalization*. One can now order a computer to one's precise specifications on the Internet, have it custom built in a Chinese factory, and delivered by an air freight carrier to your door in a matter of days. In fact, without computers and rapid worldwide communications, our present-day global marketplace would not be possible. Some authors believe that around the year 2000 we entered a new phase of this information-revolution form of globalization, what he calls "Globalization 3.0." This new phase springs from a Web-based global platform that enables multiple forms of knowledge sharing and collaboration irrespective of distance; this in turn creates a "flatter" world in which the economic playing field is getting more level between people in the developed and the developing countries. Web-enabled collaboration using wikis, open-source software, and blogs is changing the way in which knowledge is assembled, transformed, and disseminated in the information age. But ICT is not only changing commerce and industry but also transforming the power of government and big corporations to collect and assemble data on individuals.

It is also altering the way in which soldiers operate on the battlefield. Futurists predict that by the middle of the twenty-first century synergies created by the conver-

intelligence. While some, like Kurzweil, welcome such developments, others, such as Bill Joy, fear that we may not be able to control the technological genies once they are out of the bottle. Still other authors doubt that artificially intelligent machines will ever be able to master the complexities of human narrative communication. Whichever of these future projections turn out to be correct, it is certain that we will have to grapple with the social impacts and ethical challenges of twenty-first-century technologies.

SCIENCE, TECHNOLOGY, AND SOCIETY

In the modern world, technology and science often go together, with science supporting technology and technology supporting science. Although they now share a great deal in common, their goals have been historically different. In the ancient world, science, then known as natural philosophy, was viewed as an elevated activity involving pure contemplation and the value-free pursuit of knowledge, whereas technology was associated with more practical concerns and with the arts. It was not until the beginning of the modern period in the seventeenth century that there was a decisive shift to the view that scientific knowledge was valuable because it was useful to us in gaining mastery over nature. This shift was largely due to the writings of several influential philosophers such as René Descartes and Francis Bacon. Bacon's works, particularly *Novum Organon* (1620) and *New Atlantis* (1624), are notable for their contempt of traditional speculative philosophy and their emphasis on the importance of empirical methods of investigation through which the secrets of nature could be revealed by means of judicious experiments. In 1637 Descartes wrote the *Discourse on Method* in which he proclaimed that

> It is possible to attain knowledge which is very useful in life; and that, instead of speculative philosophy which is taught in the Schools, we may find a practical philosophy by means of which, knowing the force and the action of fire, water, air, stars, the heavens, and all other bodies that environ us, as distinctly as we know the different crafts of our artisans, we can in the same way employ them in all those uses to which they are adapted, and thus render ourselves masters and possessors of nature.[6]

This change in the dominant view of the nature and bases of human knowledge set the stage for the modern belief in progress, which was expressed by Bacon as the belief that "the improvement of man's mind and the improvement of his lot are one and the same thing."[8]

Despite the marriage of science and technology in the modern period, some significant differences remain between the two enterprises. Technologists primarily seek to answer the question "How?" ("How can we keep warm in the winter?" or "How can we see distant objects that are invisible to the naked eye?"). Engineers seek to design and produce useful material objects and systems that will function under all expected circumstances for the planned lifetime of the product. Science, on the other hand, may be considered as a form of systematic empirical inquiry, which seeks to describe the underlying laws governing the behavior of natural objects. Scientists primarily try to answer the questions "What?" and "Why?" ("What kind of thing is this?" and "Why does it behave the way it does?"). In the early stages of science when little

Scientists sought laws and principles that would enable them to predict and explain why things in nature behave as they do. This search produced the scientific revolution in the seventeenth and eighteenth centuries, culminating in the work of Sir Isaac Newton. But Newton's universe has been superceded by Einstein's, and later by quantum mechanics and string theory. Despite the abstract nature of contemporary physical theory, natural science continues to provide the intellectual basis for technology.

Technology needs science to predict how its objects and systems will function so that it can tell if they will work, and science supplies the predictive laws that apply to these objects and systems. However, although the laws of science are often simple to state, applying them to the complex objects of technology is often anything but simple. Sometimes the engineer must experiment with the complex objects that are the building blocks of a technology to find out what will happen. At the same time, technology makes direct and obvious contributions to the progress of science. The laboratory equipment that the scientist uses is the product of technology. The biologist would discover little without a microscope and the particle physicist even less without an accelerator. In recent years, the lines between the role of the scientist and that of the engineer or technologist have become increasingly blurred. Much of the current research agenda is dictated by the possible practical applications of new scientific knowledge, and most research is carried out by multidisciplinary teams. This merging of science and technology has led some writers, such as Bruno Latour, to speak of contemporary research as *technoscience*, a term used to draw attention to the interdisciplinary character of most contemporary research as well as the social and historical contexts in which innovation takes place.[6]

The conventional linear understanding of the relationship between science and technology holds that science discovers natural laws, technology applies scientific knowledge to practical problems, and the market selects which technologies are destined for widespread diffusion and use by society. However, this simplistic model has been replaced in recent years by a more sophisticated understanding known as the social construction of technology (SCOT) model. According to SCOT, science and technology have a symbiotic relationship, each one helping the other, while social values shape the precise forms that technological artifacts take. The new sociology of technology supports the view that technological artifacts are socially shaped, not just in their usage, but especially with respect to their design and technical content.

Social values play a crucial role in shaping technologies and in determining which of several technological options gain widespread acceptance in society. The use context of technology ultimately determines the meaning and deployment of technological innovations. Consider the Amish religious sect of central Pennsylvania and Ohio who shun the use of many modern conveniences such as the radio, televisions, video recorders, and telephones in the home because they fear that their use would destroy the rhythm of family life and cause separation to develop among members of the community. However, the Amish make compromises with modern technology, allowing flashlights, hearing aids, and electric welders for what they collectively decide to be legitimate reasons.[7] But if society's values and attitudes toward technology play a central role in determining the course that technological change takes, what attitudes

TECHNO-OPTIMISM VERSUS TECHNO-PESSIMISM

Our attitudes toward technology are complex and often ambivalent. We cannot but acknowledge and credit science and technology with delivering many wonders that have improved and extended our lives, and many people believe that improved technologies hold the solution to our problems in the twenty-first century. But many people are also disturbed by what they view as technology being out of control and see technology as a threat to our traditional ways of life, to our environment, and even to our survival as a species. These contrasting attitudes toward technology are often referred to as *techno-optimism* and *techno-pessimism*.

Techno-optimists tend to emphasize technology's benefits; they believe that science and technology are not the cause of society's current ills; they do not believe that technology needs to be controlled or regulated; and they have faith in "technological fixes" that will solve outstanding social problems. Techno-pessimists, by contrast, tend to emphasize the risks and costs of technological changes; believe that many social ills are attributable to technology; and think that technology needs to be controlled or is incapable of being controlled. They do not have faith in "technological fixes" to solve social problems, instead emphasizing moral or political solutions.[8]

While there are some extreme Luddites (those who are opposed to technological changes) and antitechnologists, the dominant view of contemporary society still seems to be a cautious form of techno-optimism. The modern idea of scientific and technological progress continues to hold sway not only for people in the developed countries but also increasingly for those in the less developed nations of the world who tend to see development largely in terms of access to more sophisticated forms of technology. However, although technological development can raise the standard of living, rapid technological and social change also brings with it social dislocation, identity confusion, a sense of disappointment, and social alienation. Part of the problem is that technology has been allowed to assume a greater and greater role in human affairs without anyone in particular being responsible for this change. Some writers see this as a problem, and others see free technological innovation as the source of prosperity and human progress.

Among the ideas that critics of technology question is the concept of progress. Throughout most of history, most societies believed strongly in tradition, and changes were presumed to be unwelcome and probably harmful. Kings sat comfortably (or uncomfortably) on their thrones, and when they were replaced through succession or conquest by other kings, quality of life changed little for the general populace. As late as 1800, life was relatively little different than what it had been in prehistoric times—most people lived in extreme poverty. Then came the steam engine, the railroad, and the automobile.

Productivity exploded in the factory and on the farm as new crop varieties and chemical fertilizers enabled fewer farmers to produce more food than ever before. As gains in productivity outstripped population growth, the industrial societies of England, Germany, and the United States grew wealthier faster than any societies in history. Telephones and railroads shrank time and space, and the factory system mass-produced goods that offered unimagined comfort and convenience to the bulk of society. Improvements in agriculture, medical advances, and public health and hygiene increased life span. In the industrialized world, progress was more than an idea; it was

In the industrialized world, however, over a century of unchallenged belief in progress was disturbed by several rude surprises. World Wars I and II demonstrated that human cruelty and brutality were still with us, only magnified by weapons capable of producing mass death. Although science and technology could put a man on the moon, the Cuban missile crisis in 1962 reminded the world that we were only a button's push away from a global nuclear war that could destroy humankind. That same year, Rachel Carson published *Silent Spring* in which she warned that pesticides such as DDT were accumulating in ever-larger amounts in species that progressed up the food chain until eagles and peregrine falcons could no longer reproduce. DDT was making their eggshells too thin to keep from cracking. It did not take a genius to realize that humans are also high on the food chain, and DDT was eventually banned. But more bad news was to follow. Mountain lakes in the northeastern United States and Europe were found to be too acidic to support fish, and the problem was traced back to acid rain, automobile emissions, and the exhaust of coal-burning electric power plants. Asbestos, our modern weapon against the age-old danger of fire, turned out to cause the lung disease mesothelioma in asbestos workers and in people living and working in asbestos-lined buildings. Radioactive by-products of nuclear power plants piled up, and no one could think of a foolproof way to keep them isolated and sealed for the thousands of years that they would be a hazard. In 1986 the Chernobyl nuclear disaster spread radioactive contamination as far away as Sweden, and the world became even more worried about the dangers of nuclear power.

Once the myth of technology as unmitigated blessing was destroyed, some people began looking for hazards posed by technology with as much fervor as had previously accompanied the search for benefits. They were not disappointed; there were heavy metals in the rivers and fish, farmland soil erosion and salinization, lead in paint and gasoline, houses built on industrial waste dumps, health problems of people processing radioactive materials, smog, ozone holes, radon, and global climate warming. Technology helped in the search for its own defects by supplying satellite photographs and instruments that could detect trace chemicals in parts per billion.

Many potential threats to human well-being have been identified, and others no doubt soon will be. Some may be false alarms that are best ignored; some may be early warnings for which action will someday have to be taken; and some may be urgent last calls for which the optimum time to respond has already passed. If technology is responsible for many of our present problems, it will likely be technology that will enable us to overcome them, sometimes in the narrow sense of finding a technological fix but more often in the wider sense that the processes of democratic decision making and economic restructuring are social technologies that we use to address and resolve social problems. Taking part in these decisions in a democratic society, however, depends on informed "technological citizens" who have attained a degree of scientific and technological literacy.

Technological citizenship is a modern moral virtue. Being a good technological citizen implies an understanding of mutual rights and responsibilities between oneself and other citizens and between citizens and the government. Among our rights as citizens are the right to receive knowledge and information about technologies and how they might affect our lives, the right to express views and opinions about the development and use of technologies, and the right to participate in decisions concerning the development and deployment of technologies that are potentially

the responsibility to educate and inform themselves about the nature of the technologies that are changing their lives and to understand the ethical and public policy dimensions of the decisions in which they claim the right to participate.

Technologies are not value neutral. In each case, there are human ends and values that stand behind and direct the technological processes. Technology itself is perceived by most people as of positive value because they understand that through technology we can increase our powers and capabilities and are therefore better able to satisfy our needs and desires. But most people also realize that technological innovations are seldom all for the good, and almost inevitably trade-offs need to be considered. A new drug may help cure a disease but may also produce undesirable side effects in some patients and may in the long run promote the spread of new and more drug-resistant forms of the disease. Car ownership may enable one to move about freely and comfortably, but it also entails loan payments, insurance payments, repairs, gasoline, smog, car accidents, global warming, and other downside effects.

Predicting how inventions and technological innovations will be used and how they will ultimately affect society is often very difficult. The history of technology is full of stories of inventors and innovators who had no idea of how their inventions and innovations would ultimately be used or the far-reaching effects that they would have on society. Johannes Gutenberg, inventor of the printing press and movable metal type, was a devout Catholic who would have been horrified to know that his invention enabled the Bible to be widely printed and so helped stimulate the Protestant Reformation. Thomas Edison apparently believed that the phonograph would be mainly used for recording people's last wills and testaments and would undoubtedly be amazed by today's tapes, CDs, and MP3 players, all of which are descended from his invention for recording sound. And who, until recently, would have thought that chlorofluorocarbons, which have been used for decades as refrigerants, would be eating away the ozone layer in the upper atmosphere? Given enough experiences of this kind, one gets the idea that every new technology has not only known and expected benefits and costs but also unknown and unforeseen benefits and costs. New technologies sometime even produce consequences exactly the opposite of what they were intended to produce, what the author Edward Tenner calls "revenge effects."[9] Powerful new technologies alter the social context in which they arise; they change the structure of our interests and values; they change the ways in which we think and work, and they may even change the nature of the communities in which we live.

Another feature of technological change is the way in which it produces winners and losers in society. If technology is a source of power over nature, it is also a means by which some people gain advantage over others. Every technological revolution has witnessed the competition among technologies and the eventual replacement of one technology or technological system by another. Think of what happened to blacksmiths when the automobile came along, or what happened to watchmakers when the quartz-electric digital watch came along, or what is today happening to bank tellers with the introduction of ATMs. In such processes of technological change, groups and individuals whose interests and livelihoods are connected to the older technology are usually the losers, and those whose interests are connected to the "next wave" of

Similar social phenomena are occurring today in the midst of the information and biotechnology revolutions and the economic phenomenon known as *globalization*. By and large, the wealthier and better-educated people in society remain largely favorably disposed toward new technologies such as computers, the Internet, gene splicing, and robots and toward the globalization of production and distribution that these technologies have made possible. Many others, however, see these developments as threatening their jobs and livelihoods, their religious beliefs, and their traditional ways of life. New technological elites are being created in each of these fields while other people are becoming newly unemployed. Such social effects of technological change bring into sharp relief the need to consider the ethical and moral dimensions of technology.

TECHNOLOGY AND ETHICS

In considering the ethical issues arising from technology, it is important to distinguish clearly between the specific products of technological development, artifacts (for example, clocks, internal combustion engines, digital computers, respirators, and nuclear bombs), and the typical uses to which people put them, or what might be termed their associated *sociotechnological practices*. The fact that a particular device or technology is available for human use does not by itself imply that we ought to adopt and use that technology, nor does it tell us how the technology should or should not be used. A gun, for instance, can be used in many ways: as a paperweight, for recreational target practice, for hunting, for personal protection, or for the commission of a crime. Although a gun has many uses, its valence lies in the social practices of use typically associated with it, which may or may not match its intended purpose. We can and do make moral judgments concerning the various sociotechnological practices associated with different products of technology. We accept some uses as morally legitimate, find others to be morally questionable or problematic, and take steps to restrict or outlaw certain other uses to which these devices may be put. In some cases, such as chemical or biological weapons whose only purpose is to produce mass death and destruction, we attempt to outlaw them entirely rather than to regulate their use. The war in Iraq that began with the U.S. invasion in March 2003 was premised on the notion that such weapons of mass destruction were present in Iraq and that, if not found and destroyed or allowed to fall into the wrong hands, could produce catastrophic results. The fact that such weapons were not found in Iraq may undermine the justification for the war, but it does not allay the concern that motivated it.

When we consider these sorts of questions about how the products of technology ought to be used, we are really asking questions about how people ought to behave or act. Questions about whether to use products of technology or how such products should be used are ethical questions; that is, they are questions concerning what we *ought* to do rather than about what we *can* do. Ethical questions related to technology are basically no different from other ethical questions that we ask about human conduct. In each case, we must attempt to determine which action or policy, from a range of alternative possible actions or policies that we might follow, is the one that we morally or ethically ought to choose. Viewed from the standpoint of technology, broadly defined, morality, ethics, and their cousin, law, are social techniques for regulating human behavior in society. They arose in human history at about the same time

cooperation based on reliable expectations that others will act as they are required to do. For instance, a simple commercial transaction in which one person buys something from someone else at a mutually agreed-on price presupposes that the buyer and seller cooperate in settling on a price and, once a price has been agreed on, in actually exchanging the goods and money that the deal requires. Such economic exchanges are regulated by social custom and, in modern societies, by a complex system of laws permitting the drawing up of contracts that legally bind individuals to the performance of the agreement terms. Other laws, such as those that prohibit theft of private property or forbid others from assault, rape, and murder, are part of a social contract that we make with one another that allows us to live together in mass societies with a reasonable degree of freedom and security.

Many people are skeptical about whether there is single, universal correct moral viewpoint. However, almost everyone believes that there is a difference between right and wrong and that most people understand that difference and can use that understanding to guide their behavior. Ethical decision making, like most other things in the modern age, is something that can be rationalized and practiced in accordance with a technique. The technique of ethical decision making consists in a conscious attempt to get a clear view of the issues, options, and arguments that present themselves in any situation that calls for ethical judgment or decision. The technique is basically this:

1. Identify all *stakeholders*—that is, all individuals whose interests might be affected by a decision.
2. Identify all possible courses of action that one might follow.
3. Review all arguments for each option, developing pros and cons in terms of their potential risks and rewards for all stakeholders.
4. Then, after having carefully worked through such deliberations, make a rational choice about which of the available options has the strongest set of moral reasons behind it.[10]

Moral reasons are those that involve ethical principles governing such notions as fairness, justice, equality, duty, obligation, responsibility, and various kinds of rights. In most ethical decisions, such reasons contend with other, nonmoral reasons for actions based on prudence or self-interest, efficiency, and economy. From the moral point of view, ethical reasons ought to always override nonmoral reasons for action when the two kinds of reasons conflict, although people do not always do what they ought to. Ethical decisions concerning the use of technologies involving judgments of value and obligation, responsibility and liability, and assessments of risk and benefit can arise at various levels: the personal level of individual behavior, the level of institutional or organizational policy, and the social level of public policy. As individuals, we are the consumers and users of the products of technology in our everyday lives; as workers or students, we belong to and participate in institutions or organizations whose policies and practices can affect our health and well-being; and as citizens, we all must be concerned about the ethical issues that we face because of modern technology.

Ethical concerns arising from technology can be divided into four kinds. The first and most basic address questions about whether and how traditional ethical values and

different ways (for example, reheat food in microwave ovens). Each new technology thus raises the implicit ethical questions: "Should we employ this new technique/technology?" and if so, "How should we employ this new technique/technology?" In many cases, such questions are answered easily. However, in many other cases, decisions about whether, how, and when to use particular technologies can raise difficult and troubling ethical issues about how our traditional ethical values and rules apply in new technological contexts.

To illustrate this kind of issue, consider how our traditional notion of privacy is being altered by modern computer and communications technologies that make it much easier to collect and analyze information about individuals. In this arena, people are asking how the traditional value that we place on privacy can be protected in the digital age. Jay Stanley and Barry Steinhardt raise the alarm concerning the increasing use of electronic-surveillance technologies by the government and giant corporations, and James Stacey Taylor argues that on balance the use of these types of technologies will make us more rather than less secure, assuming that their use is properly regulated and controlled. But as even Taylor admits, how the calculation of risks and benefits turns out will depend to a great extent on the social and political contexts in which these technologies are employed and by whom.

Traditional approaches to ethics are basically two kinds: *utilitarian* (consequential) and *deontological* (Kantian). Consequentialist reasoning in ethics involves evaluating the rightness or wrongness of actions or policies in terms of the goodness or badness of the consequences that they produce. It is often impossible to apply utilitarian, or consequentialist, reasoning to ethical problems involving technologies because it is difficult to quantify and compare the expected benefits, harms, and risks that they may produce when placed in use. One main theme of this book is that when we evaluate which new technologies to develop, which to deploy, and how to deploy them, we need to consider carefully both the benefits and costs and the opportunities and risks that the technologies entail—to the extent that we are capable of making such judgments. Often doing this sort of *cost-benefit analysis* is very difficult or extremely inaccurate because (1) manifold aspects need to be considered, (2) costs and benefits often have no common measurement scale (if they can be measured at all), and (3) we are uncertain in predicting future or long-term consequences of introducing a new technology into society. A second problem with the consequentialist approach is that it does not take into account the way in which benefits and harms are distributed and thus may give rise to allocations of social costs and benefits that are unjust. Despite these problems, consequentialist reasoning remains the dominant approach in the moral evaluation of technology.

Deontological theories in ethics emphasize not only justice but also rights and duties, which in some cases will lead to ethical judgments that would require us to follow a moral rule, honor a right, or discharge a moral duty even if doing so does not produce the greatest good. The theory of John Rawls may provide a way of combining the best elements of each approach by suggesting a way in which we can balance freedom and equality that allows each person the maximal liberty to pursue his or her own self-interest, compatible with an equal liberty on the part of others, while also requiring that deviations from equality be arranged so that they benefit the least advantaged.[11] Under this sort of view, for instance, everyone would have an equal liberty to benefit from new pharmaceutical treatments for disease, but the poorest and

these lifesaving technologies. Generally speaking, deontological considerations set limits on the possible uses of technologies and counsel us to employ our technologies only within the limits of what is ethically permissible. So, for instance, although supercomputers operated by the National Security Agency make it technologically possible for the government to monitor the billions of e-mail messages that fly around the planet each day, ethical considerations concerning civil liberties such as freedom of speech and privacy should determine what forms of electronic surveillance should be allowed.

A second kind of ethical problem arises concerning some sociotechnological practices that, although innocuous in themselves, when employed by individuals, raise serious concerns when their effects are aggregated across millions of users. There is, for instance, nothing intrinsically wrong with throwing empty bottles and cans into the trash to be carted off to the nearest landfill. But when millions of American households engage in this practice on a regular basis, we find that we are wasting recyclable resources and running rapidly out of space for new landfills. Similar sorts of *aggregation problems* arise with respect to air and water pollution, overfishing, suburban development, and many other cases in which the aggregate and cumulative effects of individual sociotechnological choices threaten the long-term well-being of all.

The current debate over global climate change due to the accumulation of greenhouse gases in Earth's atmosphere exemplifies this kind of ethical issue. The 2007 report of the Intergovernmental Panel on Climate Change (IPCC) documents the fact that "the atmospheric concentrations of carbon dioxide, methane, and nitrous oxide have increased markedly as a result of human activities since 1750 and now far exceed pre-industrial values."[12] The changes are mainly due to the burning of fossil fuels and agriculture and are thus anthropogenic—that is, caused by human activity. We are already seeing the effects of this global warming in phenomena such as the shrinking of glaciers, the defrosting of the tundra, and the reduction in Arctic ice in the summer. The IPCC predicts that unless we do something to stabilize the atmosphere, Earth may reach a "tipping point" later in this century that will dramatically alter Earth's climate and produce a significant rise in the ocean levels that will inundate many coastal areas and cause other significant environmental damage. In his popular documentary film, *An Inconvenient Truth*, former vice president Al Gore states his view that global climate change "is not a political issue; it is a moral issue, one that affects the survival of human civilization."[13] (Gore and the IPCC were awarded the 2007 Nobel Peace Prize for their work on the issue of global climate change.)

Garrett Hardin analyzed similar moral problems in his famous essay "The Tragedy of the Commons," using as his example herdsmen overgrazing common lands. In such cases, each herdsman treats the common pasture as an inexhaustible resource and seeks to maximize his own self-interest. But if every herdsman does this, the result is that the pastureland soon becomes overgrazed so that nobody can use it. Earth's atmosphere has been treated in this way by humankind throughout most all of our history, but especially since the beginning of the industrial age. We are rapidly reaching the limits of how much carbon the atmosphere can absorb without altering its geochemistry. Hardin argues that voluntary measures to limit this kind of overuse will not succeed and that the only solution available to us is greater responsibility. This

We use the term *responsibility* in both a backward-looking, or retrospective, and a forward-looking, or prospective, sense. In the retrospective sense, we think of responsibility primarily as liability for causing past harms, particularly in order to allocate blame and determine who should make amends. If we adopt this view of responsibility for global climate change, then clearly the older industrialized nations—such as England, Germany, and the United States—are responsible for the greater proportion of the greenhouse gases that have accumulated in the atmosphere and thus should bear the primary responsibility for cleaning up the mess. A second reason for allocating responsibility primarily to the older industrialized nations is that they are richer than other nations and can more easily bear the burden. On the other hand, from the prospective point of view, we still need to determine how to control future emissions of greenhouse gases, and in this case, various proposals have been made about how to allocate this responsibility for the present and future. One proposal suggests that we determine the current acceptable level of anthropogenic greenhouse gas emissions necessary to safeguard the health of the planet and then allocate shares of that amount to each country based on its population. Under this scenario, however, the older industrialized countries would still bear the greatest burden of reduction. The United States with roughly 5 percent of the world's population is responsible for emitting roughly 25 percent of the greenhouse gases, while India and China, although they are both rapidly industrializing, still are below their per capita allocations. Under most all of these scenarios, it is becomingly increasingly clear that continued delay in addressing this global problem is not an acceptable option.

A third class of ethical problems associated with technology concerns questions of *distributive justice* and social equality. New technologies generally benefit or advantage certain groups or members of society over others—namely, those who have mastery over or access to the technology first. In many cases, we think that because such advantages are earned through hard work or special knowledge they are therefore deserved. However, in other cases, we may feel that such restricted access to some technologies gives certain individuals or groups unfair advantages over others, and we seek to extend access to everyone in the society. Public libraries, for instance, were built to ensure that everyone could obtain access to books and learning. Today, we are putting computers and Internet connections into public schools for the same reason. Questions of social justice and equality of opportunity thus can be occasioned by technological innovation.

Questions of social justice are also at the heart of the debate over the current wave of globalization. Some authors believe that globalization, as it has developed over the past several decades since the advent of the information age, is a net benefit to everyone on the planet and has the potential to alleviate poverty and create prosperity worldwide. Others believe that the rules under which globalization has been conducted thus far are inherently unfair and are designed to allocate the benefits primarily to the already-rich countries and corporations at the expense of the poor and vulnerable. They argue that considerations of social justice demand that the global economic system be reformed to produce greater fairness and justice for all citizens of Earth.

A fourth and final kind of ethical question raised by technology concerns the

have the power to destroy virtually all life on Earth. Nuclear waste material from our reactors will last 10,000 years, posing a potential threat to generations as yet unborn. Issues and concerns of this type raise what are perhaps the most profound ethical questions about humankind's relationship to nature through technology. Should we continue down the course set for us by Bacon and Descartes, who advised us to seek knowledge so that we could become the masters of nature, or should we change this course toward stewardship and long-term sustainability?

Some contemporary technologies seem to open new and deeply troubling ethical issues, issues of a kind that humankind has never had to address before. The existence of nuclear weapons, for instance, forces us to "consider the global condition of human life, and the far-off future, even, the existence of the human race." The emerging technology of genetic engineering creates the prospect of our designing our own children and turning humanity itself into a kind of artifact. Some authors seem to welcome this prospect, but others believe that we are at a crossroads that requires that we relinquish the opportunity to acquire the knowledge that would enable us to create such a brave new world. Others believe that we can place reasonable limits on how biotechnology and genetic engineering will be employed on human beings that will allow some uses but prohibit others. Genetic engineering of plants and some animal species is already in widespread use, and it may already be impossible to put this particular genie back in the bottle. Hans Jonas believes that technologies such as these that give us the capability to alter nature in fundamental ways should be approached with a sense of "long-range responsibility" and, above all, a sense of humility.

ENERGY, ENVIRONMENT, AND A SUSTAINABLE FUTURE

Increasing evidence shows that our current technological society is rapidly transforming Earth's environment and probably not for the better. Hardly a day goes by that we do not hear of global environmental problems such as deforestation, species extinction, depletion of nonrenewable resources, desertification, acid rain, water pollution, ozone destruction, and atmospheric warming. In part, these problems represent the long-term and largely unforeseen effects of the Industrial Revolution, but they are also caused by the impact of human population growth and the increasing demands that consumption places on Earth's ecosystem.

There has been much discussion of the idea of a transition to systems of sustainable development, and many organizations and institutions now say they are committed to helping to bring about a more sustainable future. However, until recently the concept of sustainable development was ambiguous and ill defined, and discussions often tended to sidestep difficult questions about the real trade-offs between economic growth and environmental protection, and between the interests of the present and future generations. But a study group of the National Academy of Sciences in 1999 helped clarify the matter by defining a sustainability transition as one "that would meet the human needs for food, nurture, housing, education, and employment" for what is now predicted to be a maximum human population of about 10 billion people around

In 2000 the United Nations adopted the Millennium Development Goals (MDGs) in which the nations of the world committed themselves to the goals of eradicating extreme hunger and poverty; achieving universal primary education; promoting gender equality and empowering women; reducing child mortality; improving maternal health; combating HIV/AIDS, malaria, and other infectious diseases; and ensuring environmental sustainability.[14] According to the most recent MDG progress report, some progress has been made in some regions in meeting these goals by the target date of 2015, but much more still needs to be done. In particular, it is crucial that the richest countries honor their commitments to provide development assistance to the poor countries of the world. Particularly troubling is the continued increase of climate-warming carbon dioxide in Earth's atmosphere and the continuing migration of poor people from rural areas into already overcrowded cities. However, there are some hopeful signs. The seeds of a future sustainable society are already present, but we need to nurture them so that they continue to grow.

The global threats of the twenty-first century require social solidarity and technological innovation for their solution. These threats are different in several important ways from the threats that we faced throughout most of our previous history. First, these threats arise not mainly from the consequences of individual acts or omissions or from forces beyond humankind's control but from our own collective action. Second, they do not involve direct harms, for the most part, but rather increased risks of harm that are distributed very broadly across individuals, often without their active participation or knowledge. Third, the threats affect not only the present but also the future—often the distant, incalculable future. Fourth, they threaten not only humans but also other animals, the natural environment, and life itself. Fifth, they are also to one degree or another the result of technology; they are problems that have arisen in part because of new powers given to us by technological progress, powers that we have not always learned to use wisely and responsibly. Sixth, they not only affect single communities or even single nations but also the whole of humankind.

Our previous ethics has not prepared us to cope with such global threats. Traditional ethics has focused primarily on the moral requirements concerning individual action, on the direct dealings between persons, rather than on the remote effects of our collective action. This problem is particularly important with respect to widely distributed technologies, such as the internal combustion engine, whereby the cumulative effects of individual decisions can have a major impact on air quality even though no single individual is responsible for the smog. By and large, traditional moral norms deal with the present and near-future effects of actions of individual human beings and do not prepare us to deal with cumulative effects and statistical deaths. Traditional ethics, above all, has been anthropocentric—the entire nonhuman world has been viewed as a thing devoid of moral standing or significance except insofar as it could be bent to satisfy human purposes. We have assumed that the natural world was our enemy and that it did not require our care (for what could we possibly do to harm it really?), and nature was not regarded as an object of human responsibility.

In the past, we have attempted to fashion our ethical theories in terms of these assumptions. The traditional maxims of ethics—for example, "Love thy neighbor as thyself," "Do unto others as you would have them do unto you," and "Never treat

barriers of time, community, and species. Even more modern ethical theories such as utilitarianism and Kantian ethics do not provide particularly good guidance when it comes to the sorts of ethical concerns raised by technology. In part this is because they were designed to be used to evaluate individual actions of particular moral agents. But the sociotechnological practices that comprise our collective action are not only made up of many individual choices—such as the choice to have a child, to eat a hamburger, or to invest in a mining stock—but also the aggregation of these individual choices, plus those of organized collectivities such as corporations and governments. In most cases, the individuals, business executives, or politicians who are making the choices that add up to our collective insecurity do not intend these threats to result, and neither they nor we consequently feel any sense of responsibility for them.

Although individuals view themselves as moral agents and consider themselves bearers of responsibility in all the roles in which they participate, the collectivities to which we belong do not. All the threats that we face are in part the result of this diffusion of responsibility. How then should we, the citizens of Earth, be responding to these environmental questions? Do people in richer countries have any responsibility to help those in poorer ones? Do we, in general, have any responsibilities to future generations concerning the long-term social and environmental effects of our present economic, lifestyle, and political choices? The notion of responsibility that we need to cultivate is not the backward-looking notion of responsibility as liability, which seeks to allocate blame for past harms, but the forward-looking sense of responsibility in which each of us and every organization and institution "takes responsibility" for doing our part to combat social injustice and to protect environmental quality for future generations of humans and the nonhuman species with whom we share this planet. This notion of social responsibility, although it is voluntary and discretionary, places real demands on us as individuals and members of communities and requires that we think carefully about the decisions and choices that we make.

All too often, decisions that involve complex political choices involving technologies are left to the discretion of elites (for example, scientists, engineers, policy "wonks," and corporate and government officials) even though the consequences of their decisions will usually affect the interests of others who are not elites. The other interested but often silent parties are sometime called *stakeholders*. We are all stakeholders in decisions concerning technology, but not infrequently the scientific, political, or corporate elites make decisions about these questions in ways that primarily benefit themselves at the expense of other stakeholders. It is often relatively easy for elites to "manufacture consent" for policies that they prefer by selectively sharing information about the possible risks and benefits of a particular technology policy with other stakeholders whose interests might be adversely affected by it.[15] For instance, in the 1950s U.S. soldiers were ordered to witness nuclear explosions and were told that there was no risk of harm due to radiation. In fact, there was a risk, and years later many of the soldiers who participated in these tests began developing lethal cancers. More recently, automobile companies, such as General Motors, conspired with giant oil companies and corrupt officials to "kill" a prototype electric car, the EV-1, despite consumer interest in an economical and nonpolluting alternative to petroleum-based personal transportation.[16]

and Health Administration (OSHA), which are mandated to act as watchdogs and look out for the interests of the public and to prevent people from being exposed to unnecessary or unreasonable risks without their consent. However, the operations of these very governmental agencies have often become politicized, and key officials appointed to run these agencies sometimes represent corporate interests rather than the public interest.

Given the phenomena of regulatory capture by corporate interests, a more reliable line of defense is the hundreds of nongovernmental organizations (NGOs), such as Common Cause, Greenpeace, or the International Center for Technology Assessment, that conduct independent research, educate the public, and lobby decision makers to enforce and protect the stakeholder interests that they are supposed to represent. Such public-interest groups and the social movements that they represent play an important role in politics and provide a means, in addition to the ballot, by which ordinary citizens can participate in large-scale decisions that many affect their lives for good or for ill.

However, none of these advocacy groups can be effective without the support of an informed and attentive citizenry. In democratic societies, individuals and groups are given the right to inform themselves on the issues, associate with others having similar or common interests, and participate in the political discussions that will determine which laws and policies will be enacted. If we fail as individuals to exercise these rights—that is, if we shirk our responsibilities as technological citizens—it is likely that others will end up making these decisions for us, and when they do, they may not always have our best interests at heart or in mind. If we accept the responsibility to educate ourselves about the issues and to participate in the public conversations about them, then we will have some voice in how things will be decided and some control over the future directions that our technological society will take. In the last analysis, there is no way for us to escape this responsibility, living as we do at the cusp of the third millennium, for we are now all the children of invention.

Morton Winston
April 2012

NOTES

1. For a rather long but still incomplete list of some of humankind's most significant inventions, see the Time Line of Significant Technological Innovations on the inside front and back covers.
2. Compare this definition to that found in Rudi Volti, *Society and Technological Change*, 2nd ed. (New York: St. Martin's Press, 1992), in which technology is defined as "a system based on the application of knowledge, manifested in physical objects and organizational forms, for the attainment of specific goals" (p. 6).
3. The term *valence* is also used to describe the way in which tools and technological systems have "a tendency to interaction in similar situations in identifiable and predictable ways." The terms *end* and *focal function* refer to the purpose in the mind of the designer of the artifact. See Corlann Gee Bush, "Women and the Assessment of Technology: To Think, to Be; to Unthink, to Free," in *Machina ex Dea*, ed. Joan Rothschild (New York: Teachers College Press, 1983), 151.
4. René Descartes, "Discourse on the Method of Rightly Conducting the Reason and Seeking for Truth in the Sciences" (1637), *The Philosophical Works of Descartes*, Vol. I, trans.

E. S. Haldane and G. R. T. Ross (Cambridge, England: Cambridge University Press, 1970), 119.

5. Francis Bacon, "Thoughts and Conclusions," in *The Philosophy of Francis Bacon*, ed. B. Farrington (Chicago: University of Chicago Press, 1964), 93.
6. See Bruno Latour, *Science in Action* (Cambridge, MA: Harvard University Press, 1987).
7. See Donald B. Kraybill, *The Riddle of Amish Culture* (Baltimore: Johns Hopkins University Press, 1989), especially Chapter 7.
8. The terms *techno-optimism* and *techno-pessimism* were suggested by the discussion of pessimism and optimism about technology found in Mary Tiles and Hans Oberdiek, *Living in a Technological Culture: Human Tools and Human Values* (New York: Routledge, 1995), 14–31.
9. See Edward Tenner, *Why Things Bite Back: Technology and the Revenge of Unintended Consequences* (New York: Random House, 1996).
10. For more on ethical decision making, see C. E. Harris, Jr., *Applying Moral Theories*, 5th ed. (Belmont CA: Wadsworth, 2006).
11. John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971).
12. Intergovernmental Panel on Climate Change, "Contribution of Working Group I to the Fourth Assessment Report," February 5, 2007; available online at http://www.ipcc.ch. Accessed July 7, 2007.
13. Davis Guggenheim, *An Inconvenient Truth*, Paramount Pictures, 2006.
14. United Nations, "Millennium Development Goals"; available online at http://www.un.org/millenniumgoals. Accessed July 9, 2007.
15. The idea of manufacturing consent is based on the work of Noam Chomsky, See especially, Noam Chomsky, "The Manufacture of Consent," in *The Chomsky Reader*, ed. J. Peck (New York: Pantheon Books, 1987), 121–136.
16. Chris Paine, *Who Killed the Electric Car?* Sony Pictures, 2006. See also the sequel, *Revenge of the Electric Car* (2012).

PART ONE

Perspectives on Technology

1.1 Historical Perspectives

1.2 Social/Political Perspectives

1.3 Ethical Perspectives

1.1 HISTORICAL PERSPECTIVES

1.1.1 SOMETHING FOR NOTHING

JAMES BURKE

Science historian, author, and television producer, James Burke, is perhaps best known for his 1978 BBC television series "Connections," which explored a wide range of technologies that had links with later innovations. Contrary to the popular view, there are very few actual "Eureka" moments when some great discovery pops into an inventor's mind. The path followed more often resembles the seemingly random walk that Burke describes.

In this reading, taken from his book *The Pinball Effect: How Renaissance Water Gardens Made the Carburetor Possible–and Other Journeys through Knowledge,* Burke begins his tale with the solution developed in order to pump water from deeper and deeper mines in England, and progresses through a number of seemingly unrelated discoveries, eventually bringing three inventions together to show how they were essential to the creation of something totally unrelated, the space shuttle.

The survivability of the space shuttle crews depended on the ceramic tiles which protected the vehicle from very high entry-level temperatures generated by air friction, the gyroscope which permitted accurate navigation of the shuttle, and the fuel cells which provided the reliable power essential for sustained flights.

Even though the total of 135 flights by U.S. space shuttles was far below their projected design life, the fact that these remarkable vehicles were conceived, constructed, and flown before the age of the high-speed computers on which much technology relies today is amazing. Their existence depended on many developments that were essential to the operation of the total system. The space shuttle did not spring from the mind of a single person at one moment of inspiration but grew from many discoveries and inventions that were creatively recombined. This selection then illustrates the often surprising synergies among seemingly disparate technologies.

FOCUS QUESTIONS

1. In what way did religious dogma impede scientific research in regard to barometers and pose a threat to scientists and conclusions they might reach? Is there any evidence that this has happened in other times in human history?
2. What changes during the time periods discussed by Burke provided the impetuses that led to further discoveries?
3. Along the road of exploration described in this selection, what were some of the intermediate discoveries made, and how did they benefit society?
4. Given the current reluctance by many politicians to impose taxes or increase debt to pay for public-works projects to support our infrastructure, what might have been the outcome if a similar philosophy had prevailed during the 1800s?

Source: "Something for Nothing", from *The Pinball Effect* by James Burke, pp. 133–147. Copyright © 1996 by London Writers Ltd. By permission of Little, Brown and Company. All rights reserved.

KEYWORDS

alternating current, atmospheric pressure, barometer, battery, ceramic tile, cholera, compulsory vaccination, direct current, fuel cell, gyrocompass, gyroscope, inertial navigation/control system, John McAdam, Meteorological Office, muscle movement, pneumatic, suction pump, vacuum, vitreous enamel, Westinghouse air brake

This historical journey shows how a single discovery, over three centuries ago, was to trigger the eventual appearance of three very different, rather ordinary modern inventions: heat-resistant ceramics, the gyroscope and the fuel cell. These three inventions made possible a fourth one, which is one of the most extraordinary vehicles on the planet.

The story starts in mid-seventeenth-century Europe. At the time, a mining boom was in progress—or it would have been if the miners hadn't come across a serious problem. For some strange reason, the suction pump (the age-old system for draining the mines of the water that plagued a miner's life) wouldn't lift water more than thirty feet or so. At this point the miners had to introduce another suction pump, which would lift water only the next thirty feet, and so on. Several hundred feet below ground, this system became a serious waste of time and money. So concern was expressed to the authorities and thinkers of the day.

In 1639 the Genoese governor Giovanni Baliani told his friend, the scientist Galileo Galilei,that he'd had the same trouble trying to siphon water along a pipe and over a hill. Galileo gave the job of finding the answer to his pupil Evangelista Torricelli, who told various Roman friends. In 1641 one of them, Gasparo Berti, rigged up a device consisting of a tube full of water, with its lower, open end immersed in a flask of water and its upper end closed with a tap. When the tap was opened, the water in the tube ran out into the flask until the level in the tube had fallen to a point about thirty feet above the water level in the flask below, whereupon it stopped.

Torricelli tried the same trick with mercury; since mercury is denser than water, he needed much less of it and could experiment on a more manageable scale. It behaved the same way as the water. Torricelli also told Marin Mersenne, a cleric (and the man with the Continent's biggest address book of amateur scientists), who in turn told the whole of Europe. The result was that on September 19, 1648, just outside the central French town of Clermont-Ferrand, a bunch of local worthies and the brother-in-law of famous French mathematician Blaise Pascal (who was a pal of Mersenne's) climbed the nearby Puy de Dome mountain, with a tube of mercury upended in a dish of mercury. At the bottom of the hill they repeated Torricelli and Berti's test. The mercury in the tube fell a certain distance and then stopped. As they climbed the mountain, the level of mercury in the tube gradually fell farther. And as they descended, it rose again. This clearly meant that the atmospheric air pressure, pushing down on the surface of the mercury in the dish, was supporting the column of mercury in the tube. That was why the miners' suction pumps wouldn't work vertically more than thirty feet. Above that height, air pressure was insufficient to support the column of water.

The miners' problem was solved, but another was created: as the experimenters had climbed the hill and some of the mercury had run out of the inverted tube, the falling mercury column had left a small empty space at the top of the tube. This space could not be air (since none had bubbled up through the mercury), so it had to be a vacuum. According to the pope, this was impossible. A vacuum would mean totally empty space and—according to the church—God was supposed to fill the whole of space. So the existence of the vacuum was a dangerous thing to mention in a Catholic country. Which was why research from then on moved north of the Alps, into Protestant lands.

But in spite of the relative freedom of investigation there, it would take two hundred years and

a major disaster in Russia for the world to reap the full benefits of the original mountain climbers' airless discovery.

It would be the first in a sequence of events triggered by the vacuum that would lead to one of the three inventions in the modern world with which this tale ends.

Since the vacuum in the top of the mercury tube went up and down according to the altitude, obviously the air pressure on the mercury in the dish changed with height. It wasn't long before experimental amateurs in the English Royal Society noticed that this vacuum also rose and fell with changes in the weather. In consequence, by 1708 few gentlemen of quality were without the hi-tech wonder-instrument that this discovery had generated. The instrument became known as the barometer (colloquially, the weather glass), and any experimenter worth his sodium chloride had to have one. By 1720 it was known that the mercury would stand thirty-one inches up the tube in very dry weather and only twenty-eight inches in storms. Ladies were advised to have a weather glass nearby to consult when they dressed, so as to "accommodate their habit to the weather."

By early in the nineteenth century, weather data had been recorded for long enough to show that, on average, it rained every other day, more frequently at night, most in autumn and least in winter, most in October and least in February, and that one year in five was very dry, and one in ten very wet. Not surprisingly, the reason for this interest in the weather was financial. Every year storms were sending thousands of pounds' worth of cargo (and, on average, two thousand crew members) to the bottom of the sea.

In 1835 a researcher at the Philadelphia Franklin Institute called James P. Espy (about whom remarkably little is known) started collecting data from up and down the east coast of America, as far north as Canada and along the coastline of the Gulf of Mexico. He sent out hundreds of circulars to sailors, to lighthouse keepers and to six hundred army posts. After four years he had received enough data to draw over a thousand storm-track maps, from which he was able to deduce that storms were either round or oblong. By 1839 Espy was world famous, lecturing in the United Kingdom and Europe, known as the author of *The Philosophy of Storms*. Interest in the weather was now like the barometer in good weather. It was rising.

As early as 1843 England's Committee on Shipwrecks was recommending that every vessel carry a barometer. On land the first weather map was produced by a British navy captain Glaisher in June 1849 for the *Daily News* and sold to the public for a penny. Naturally enough, naval interest in things meteorological was high both in Europe and in the United States. In Washington, D.C., Matthew Maury, a young navy lieutenant, amassed a mountain of data on winds and currents copied from captains' logs over a period of nine years. In the end he had the equivalent of a million days' worth of observations. This was enough to persuade everybody to get together at the world's first weather conference in Brussels in 1853. As a result of this conference, nothing happened.

Then, a year later, on November 14, 1854, the Crimean War Allied fleet found itself anchored off the Russian port of Balaklava. The fleet consisted primarily of supply ships like the HMS *Prince*, the greatest steamship of the time, but there were warships, including the *Henri IV*, pride of the French navy. During the afternoon, violent rainstorms began to lash Balaklava, and by nightfall a full-scale hurricane was blowing. By morning the entire fleet had been sunk. The *Prince* alone took seven thousand tons of medical supplies, boots and forty thousand winter greatcoats to the bottom. As a result of the losses, Allied troops ashore then had to live through a terrible winter without supplies. The appalling conditions under which they survived and the consequential high death rate hit the headlines everywhere, triggering, among other things, Florence Nightingale's famous mission, which would lead to the total reorganization of military, and then civilian, hospitals all over the world.

Within days of the storm, the French emperor Napoleon had ordered an investigation. His academician Urbain Leverrier collected 250 sets of weather data from every European observatory for the period November 11 to 16. His report, issued at the end of January 1859, shocked

everybody. The Crimean fleet could have been saved because it could have had at least twenty-four hours' warning. The day after receiving Leverrier's report, Napoleon ordered a telegraphic storm-warming system to be set up all over France. In the United Kingdom the response was just as swift. There the board of trade established the Meteorological Office, and by the summer of 1860 it was receiving daily data from fifteen stations around the British Isles, as well as exchanging telegraphed weather reports, via Paris, with a dozen places throughout Europe.

Interest centered on the weather in and around Oxford Street and the London Embankment. This was where major new road works had been under construction through the second quarter of the century. By 1850 London had a rapidly rising population, and evidence of its mobility around the city manifested itself in the form of the twenty-four tons of manure being collected daily from Oxford and Regent Streets alone. With true Victorian mania for statistics, the planners reckoned this amount represented eight thousand daily horse droppings. In addition to the mud and loose stones with which the streets had previously been paved, cleanup costs and inconvenience were considerable. A gold medal worth thirty guineas was offered for a new street-cleaning technique. In 1844 the *Times* of London reported on progress so far: "Forty to fifty men and boys … employed during the day in the removal of excrementitious matter … [proves] that the experiment of washing the streets … has been … found successful."

In 1847 the local commissioners for roads began to introduce a new surfacing technique, invented by John McAdam, which consisted of firm roadbed foundations topped by granite chips or paving stones. At last, the streets could be washed down easily. But beginning in 1855, this would cause nothing but complications for the engineers digging London's massive new sewerage systems. Road-washing and stormwater runoff from the new impermeable road surface represented a major factor in estimating how big sewers ought to be. This calculation was accomplished with the aid of meteorological data and by referring to information collected in Paris and Lyons, where the French already had rain gauges set on nearby watersheds to warn of impending flooding.

Water and sewage discharge was analyzed as carefully as is electric power use today. In the morning, water would be emptied from baths. Before noon, kitchen washbasins, water closets and housemaids' sinks would be "freely discharging various foul matters." In the afternoon and evening, kitchen sinks would contribute most. The average sewage flow for London was reckoned to be 15,208,083 cubic feet per day. There was an average 8 percent fluctuation in peak hourly flow.

These statistics were made more complicated by the weather. Rainfall records indicated twenty-one days a year when more than a quarter inch fell. Heavy rains could cause disaster in many old sewers if the average sewage flow were less than 2½ feet per second for more than a few days. And severe storm flow could be as much as six times that of the flow in dry weather. Pipe leakage was an additional major problem. Nonetheless, the engineers pressed ahead. When the new sewers were finished in 1871, they stretched for eighty-three miles beneath the streets of London to the suburbs in the east (downriver, where the sewage would be dumped), and the system drained over a hundred square miles. The construction had cost the staggering sum of £4.6 million sterling.

Now, people spend that kind of money only when they're desperate. The cause of the desperation was an epidemic that visited Europe three times between 1831 and 1866, and in England alone killed over a hundred thousand people. It was cholera, and it created total panic because it struck both high and low throughout the land. Nobody was safe, and nobody knew how to stop it. All they knew was that it hit the poor and dirty harder than the rich and clean. And there were more poor and dirty than ever before. In the overcrowded and exploding industrial cities, immigrant factory labor from the country villages packed into tenements, living in unspeakable conditions of filth, starvation and disease. One in two children died before the age of five, and only one in six adults made it to the age of fifty. Between 1831 and 1841 the average death rate from

starvation and disease rose by a third. The country was already close to anarchy when cholera struck.

The Victorian middle class knew it was on the verge of revolution unless something drastic were done. So they built sewers, started Associations for the Diffusion of Sanitary Knowledge and sent their children out of the cities to schools where sports and muscular Christianity and cold showers might protect against the epidemic. They fumigated and whitewashed the slums, and Methodists held anticholera pray-ins. The marriage rate rose. Local authorities offered anticholera candies to postmen, and national compulsory vaccination was introduced. By 1875 Parliament had enacted no fewer than twenty-nine pieces of sanitary legislation, ranging from the Public Health Act to nuisance removal measures. Most important of all, the government removed the excise tax on soap; so between 1841 and 1861, its use doubled. Cleanliness was now next to godliness. Some 140,000 tracts were distributed on the "power of soap and water."

All this was music to the ears of pottery makers, who saw the greatest-ever new market opportunity emerging from the filth, death and degradation. Potters who had spent generations making earthenware Toby jugs now made a fortune glazing sewer pipes. The Doulton firm, among others, became a household name with its new ceramic kitchen sinks and pitcher-and-bowl china toilet sets. Plumbing soon linked houses to the new sewers, and then everybody wanted a water closet. Being Victorians, they wanted them respectable, so the manufacturers made them in decorous white vitreous china, painted with designs called Magnolia, Wild Rose, Unitas, Renaissance, Baronial, Olive Green Chicago and Morning Glory. For those more concerned with their sanitary efficacy, names included Directo, Precipitas, Inodoro, Rapide and Deluge. The first properly modern flush lavatory was sold in 1884 by George Jennings (it won prizes and introduced the oval seat design), under the brand name of Pedestal Vase.

The first of this chapter's three modern inventions finally emerged when Victorian hypochondria brought new Industrial Revolution techniques in cast iron together with the kind of specialist glassworking skills which the new chemical industries were demanding, especially in Germany. Chemical producers needed containers that would be impervious to the effects of corrosive liquids. For this reason, the Germans (the Continent's best chemists) were particularly good at enameling. Since the 1840s they had been treating cast-iron cooking pots, and then all kinds of kitchen-ware, with the blue enamel which would become so familiar to Europeans and to American frontier settlers by the end of the century.

The new boom in sanitary ware boosted the ceramics industry as never before. By 1904 no fewer than seventy German companies were producing cast-iron baths covered with vitreous enamel which was little different from the kind found in bathrooms today. And the process of manufacture had been refined to a point where the most complex of shapes could be covered using liquid enamel (made by producing the ceramic material in high-temperature furnaces, then running the molten mixture in cold water to splinter it and grinding the splinters down to fine dust). At this point, the dust was remelted into liquid at 1,000°F (542°C), and the hot cast-iron ware was either sprayed with it or dipped into it. A final dusting with glass powder gave the necessary smooth finish.

This ceramic technology was the first of the three modern inventions which come together at the end of this chapter and make possible one of the most extraordinary feats of exploration in history.

The second of these three modern inventions came about as a result of the rapid seventeenth-century development of the vacuum as an experimental tool. Creating a vacuum involved sucking air out of the instrument which was to contain it. The first step toward being able to do that was taken in 1654 by Otto von Guericke, mayor of the German city of Magdeburg, when he designed an air pump. In London a Royal Society experimenter, Robert Boyle, used a modified version of the pump that would either increase or decrease the air pressure in a vessel. Robert Boyle placed a barometer in such a vessel and saw that as air pressure rose, the mercury in a barometer caused trapped air to become squashed. When pressure was released, the trapped air would expand and

push back on the mercury. Air, apparently, was flexible, and could be compressed.

Various schemes were suggested to utilize this exciting new property. In 1653 the Frenchman Denis Papin designed the first pneumatic tube for delivering parcels. In 1785 a Scotsman suggested raising sunken ships with compressed air balloons, and after the great storm of 1854 some of the ships sunk in Balaklava were raised this way. The first professional diving also began.

It was news about one particular use of compressed air that would help open up the American West. In 1857 work had begun on the first tunnel under the French Alps, which at the time were Italian, as Savoy still belonged to the king of Sardinia. The plan was to use the tunnel to unite the railway systems of Italy and northern Europe. This would cut the journey time between Paris and Turin to eighteen hours, radically improve freight delivery between the Mediterranean and the north, open up the Italian Riviera to tourism and shorten the British return journey from India by several weeks.

Elsewhere, this first Alpine tunneling work (the tunnel under Mont Cenis, running from Mondane in France to Bardonecchia in Italy) raised even more ambitious hopes. One day in Pittsburgh a young American bought a magazine subscription because he liked the looks of the girl selling them. The first copy he read carried a story about Mont Cenis and the compressed-air rock drills being used there for the first time (and which were increasing tunneling progress, over hand-drilling techniques, twentyfold).

The young magazine reader was George Westinghouse, and like many Americans he was well aware that railroad travel was risky business. Every month another horror story hit the headlines as overloaded freight or passenger trains with inadequate brakes crashed on steep gradients, with horrifying loss of life. At the time, every car carried a brakeman, and in an emergency, his operation of the manual braking system was supposed to coincide with those being applied on all the other cars. Casey Jones was the most famous example of what could happen if things went wrong, and cars jackknifed into each other with disastrous consequences. As one contemporary columnist wrote: "Another railway accident (so-called) on the Erie Road. Scores of people smashed, burned to death or maimed for life. We shall never travel safely till some pious, wealthy and much-beloved railway director has been hanged for murder."

Westinghouse realized that the answer lay in compressed air. By 1869 he had a patent for the Westinghouse air brake, which has been used in some form ever since. In its automatic version, compressed air in a pipe running the length of the train (and backed up by compressed air reservoirs under each car) held back a piston. If the air pressure were released (deliberately or by accident), the piston closed, putting on the brakes. The system was a (non-) smash hit. By 1876, 15,569 locomotives and 14,055 cars were equipped, and the system could stop a 30-mph, 103-ton train in 500 feet, three times shorter than the old way.

Through his invention, Westinghouse became involved with a Croat genius who wore a new red-and-black tie each week, used handkerchiefs once and threw them away, hated billiard balls and lived in a hotel room full of pigeons. (But more about him later.) The Westinghouse brake and the now-minimal risk of crashing meant trains could run faster and more frequently. It was important to be able to control traffic movement better, so in 1880 Westinghouse began buying up patents in signaling systems. By 1882 he had set up the Union Signal and Switching Company in Pittsburgh. At first he tried controlling the signals along the track by compressed air, but the pipes froze in winter. Electricity was clearly going to be the answer, and his idea was to send an electric current that would operate a valve to release compressed air that would work the signal.

Westinghouse started a business in power transmission, where the problem at the time was direct current. In 1871 electricity was produced by rotating a ring carrying copper-wire coils between the poles of a magnet made of an iron bar, itself coiled with wire. In this self-exciting, "dynamic" system (from which the name "dynamo" comes), as the coil turns, it sets up current in the wire round the iron bar, making it a magnet. This in turn sets up magnetic lines of force. As each side of the spinning ring cuts the

magnet's lines of force going in opposite directions, its coils generate current flowing first in one direction, then in another. This current was made unidirectional by using a device (a commutator) with separate contacts picking up and passing on first one current, then the other. Direct, unidirectional current was the result.

The trouble was that low-voltage, direct current (needed to power lights, for instance) could be sent only about a mile, down thick, expensive copper wire, and a generator was required for every mile the current went. However, without a commutator, the system produced alternating current. This could be sent down thin, cheap wire at very high voltage. If, at the source, a transformer stepped up the voltage (the "pressure" of the current) so it would go long distances, and at the receiving end other transformers stepped the voltage down again for local use, alternating current could free electricity from the constraints of geography. Since electricity could also be made by spinning the magnets instead of the wire coils (which, at high speed, often flew apart), what was needed was a high-speed power source to spin the magnets fast enough to produce high voltage.

In 1888 Westinghouse met that pigeon-loving Croat, Nikola Tesla, who had invented just such a system using alternating current. His idea was brilliant. If the current were sent in separate sequential, bursts—into two or more sets of coils wound on iron—a series of currents out of phase with each other would be set up. As each burst of current happened, around the ring the bursts would create a sequential series of magnetic fields. The magnetic fields would, in effect, rotate—and this could be used to make a magnet spin at virtually unlimited speed. By 1895 Westinghouse had designed a complete generating system around Tesla's idea, winning the contract for the giant new power station at Niagara Falls.

Tesla's motor would also turn out to be just what battleship captains were looking for, as the American navy entered World War I. Thanks to recent advances in metallurgical techniques, new, giant all-steel, armorplated heavyweight battleships had been introduced by the British, who in 1906 launched the HMS *Dreadnought*. The "dreadnoughts" were also the first "all-gun" turbine-powered warships, and their firepower was extraordinary. New gun-making techniques, using wound steel wire, permitted the construction of longer barrels needed to match the accuracy with which new, metal-jacketed shells could now be fired, because of new, slow-burning explosives. With piston-operating shock absorbers preventing recoil, the guns could also be quickly breechloaded from behind. And because of new, smokeless powder, there was neither flash nor smoke to give away the ship's position or obscure the gunner's aim.

By 1917 the American dreadnoughts USS *Delaware*, *Utah*, *South Dakota*, *Florida*, *Arkansas*, and *Wyoming*, each displacing 226,000 tons, carried fifteen-inch guns (mounted fore and aft) that could throw a two-thousand-pound shell nearly fourteen miles. Accurate fire control now became an urgent necessity, because at the extreme range these new guns permitted, the slightest pointing error meant a miss. The solution to this problem was, fortunately, already being developed by Elmer Sperry, a New York electrical manufacturer who had created the hi-tech version of a children's toy and had failed to interest Barnum and Bailey's Circus in equipping a high-wire wheelbarrow with one. By 1912 Sperry was in the process of installing several multi-ton versions of his new toy on ships, including the Italian luxury liner *Conte di Savoia*, which in consequence became one of the smoothest-sailing ships afloat.

His amazing new device was called a gyroscope. In principle much like a child's top, the gyroscope consists of a spinning wheel mounted on three-way gimbals. The effect of inertia means that as long as the wheel spins, it will always point the way it was originally set. Given the problems that the metal construction of the new, all-steel ships created for any magnetic compass on board, a gyrocompass was a godsend. And thanks to Tesla's electric motor, there was now an electrodynamic means to keep the gyro wheel spinning indefinitely.

Elmer Sperry installed in ships a master gyrocompass that controlled all other compasses on

board. He also designed equipment, using the gyrocompass readings, that would track the relative location of a target, would show the ship's own position at any time and would provide the gunners with an arrow showing how they should point their guns. The ability of the gyro to sense the ship's movement on a rolling sea and send compensatory signals so as to keep the guns correctly aligned also meant the ships were now a stable platform, no matter what the sea conditions. Gyro-controlled guns on the USS *South Dakota* shot down every aircraft during an attack in heavy seas.

Sperry's gyrocompass is the second of this chapter's three modern inventions, originally derived from the discovery of the vacuum.

The third of these modern inventions emerged because the seventeenth-century discovery of the vacuum got people excited about air: what it is and, above all, why it is essential to life. Respiration of all kinds came under the scrutiny of English Royal Society amateur scientists. One of them was a Cambridge-educated clergyman living in Teddington, on the River Thames, just outside London. His name was Stephen Hales, and in 1709 he took up the perpetual curacy of Saint Mary's Church. This job-for-life meant that he could devote all his spare time to his two obsessions: how plants breathe and what makes muscles move. Hales was convinced that the answer to both conundrums lay in the movement of fluids in plant and animal bodies, so he started by looking at why sap rises.

In his Teddington greenhouse, over a period of ten years he stuck glass tubes into plants at various points and watched how far up the tube plant sap would rise at various times, and in various weathers, day and night, year-round. He measured plant-leaf areas and stalk diameters; and then he sealed them in their pots, fed them measured amounts of water and weighed them to see how much water loss occurred through transpiration and evaporation. He grafted onto a plant a tube containing water sitting on top of mercury, measuring the mercury rise as the plant took up the water. What most excited Hales was the force of the sap. Sticking a freshly cut vine stump into one of his glass tubes, he was amazed to see the sap eventually rise to twenty-one feet up the tube. What, he asked himself, was the force driving the sap? And was it the same force that moved blood around the body?

After unspeakable experiments on the various arteries and veins of deer, oxen, mares, dogs, cats, rodents and assorted livestock, Hales worked out that if a glass tube were to be fixed to a many carotid artery, and the artery were severed, the blood would go 7½ feet up a tube. Hales thought the heart's left ventricle measured fifteen square inches, so a 7½-foot rise indicated that the heart was delivering 51½ pounds of force. In the long run, Hales came up with the first basic book, used by doctors ever since, on blood pressure and how to measure it. Regrettably, he found he was mistaken about how muscles are moved. There just wasn't enough blood pressure in the capillaries going into the muscles to produce the kind of force that muscles exert.

However, not long afterward, at the northern Italian university of Bologna, Luigi Galvani, the Professor of Anatomy, went looking for the same force. Now, in 1780, most people thought if the force wasn't blood pressure it was at least some kind of invisible fluid pumping down from the brain and causing the muscle fibers to become "irritable," so that they contracted. If this were the case, was the force some kind of electricity, the latest wonder-force? Electric shocks seemed to cause everything from flashes before the eyes to nosebleeds, convulsions and jumping in the air. Galvani believed the "fluid-along-the-nerves" theories were wrong, because if a nerve were cut, electricity still flowed through the local area. But in the newly discovered torpedo electric fish, only certain parts of the fish gave a shock. So the force, whatever it was and however it acted, could be confined. And if it could be confined, it could be directed (to muscles, for instance).

Galvani started on frogs and through long, hot, thundery Bologna summer days (it can get really long and hot and thundery in Bologna), he subjected pairs of frog legs, attached to their spine (and cut to reveal the main nerves), to various shocks. He noted that they jerked when his

electricity machine sparked and when there was lightning around, but only if his scalpel were touching the frog nerves at the same time. This was inexplicable.

As the days drew soporifically on, the lightning flashes became infrequent enough for things to get really boring, so, almost off-handedly, Galvani began to scrape the copper hook on which his frog legs were hanging against the iron of the stand on which the experiment was mounted. The legs contracted. And they did so every time he scraped. In 1791 a long, rambling and confused Latin monograph from Galvani announced he had found the source of the force. It was in the frog. Animals obviously made electricity. What else could account for the phenomenon?

A contemporary Italian electricity freak, Alessandro Volta (at yet another one of those very busy crossroads on the web, thanks to the fundamental role he plays in history) had the answer. The force was the frog all right, but only because it was playing the role of anything salty and wet (which all organic cells are), so the two different metals Galvani was using (copper and iron) were reacting through this wet, salty frog and producing electricity. Volta proved this by repeating it by another means. He connected a series of cups, filled with a brine solution, via links made alternately of zinc and copper, and got a shock from the final wire link. And the greater the number of cups, the bigger the shock. To make this arrangement more portable for experimenters, Volta stacked discs of zinc and copper, sandwiched between discs of cardboard soaked in brine. Bringing together the ends of wires attached at each end of this pile of discs (Volta called it his "pile") caused a spark to happen. Thanks to Galvani's wet frogs, Volta had produced the world's first battery.

It was not for another fifty years that people actually began to understand what was happening in the salty liquid to produce the electric current: the molecules in the salty solution were adding electrons to one of the metals (the cathode) and taking electrons from the other (the anode). This interaction caused a current to flow in the link (the salty liquid) between the metals. But in doing so, the passage of current gradually dissolved one of the metals. In 1889 a couple of Germans called Mond and Langer found a way to avoid dissolving the metal, making possible an efficient battery. Ironically, Tesla's dynamo killed the idea, and it was not revived again till World War II.

Today, the device has come into its own—thanks to our ability to liquefy gases—in a form that replaces discs and wires with porous electrodes placed in the solution. From liquid-gas sources, one electrode is provided with oxygen, and the other with hydrogen. The hydrogen diffuses through its electrode, which absorbs the hydrogen atoms. The atoms then react with the solution to make water and, in doing so, give up electrons to the electrode, in the form of current. The water and the current are passed to the other half of the setup, where the electrode is busy absorbing oxygen. This electrode reacts with the oxygen (and the water from the other electrode) to create hydroxyl atoms that then migrate through the water to the hydrogen electrode, completing the circuit. The product of this system (called a fuel cell) is electricity and water. This is the third and final modern invention derived from the original discovery of the vacuum.

And, together with the other two inventions (the gyroscope and ceramic tiles), it helps to make possible a fourth invention. First, this needs a constant and autonomous supply of electricity and water provided by a fuel cell. Second, it has to have a gyroscope that will give direction and position unaffected by magnetic or gravitational influences. Third, it needs ceramic protection against temperatures as high as 3,000°F (1,662°C).

All three inventions work together. The fuel cell powers all the onboard instruments; the gyroscopic inertial navigation-and-control system enables the vehicle to orbit the Earth before its pinpoint, one-attempt-only landing; and the crew survives the fiery furnace of re-entry thanks to the 34,000 ceramic thermal tiles covering the belly of the fourth technological marvel: the space shuttle.

1.1.2 INDUSTRIAL SOCIETY AND TECHNOLOGICAL SYSTEMS

RUTH SCHWARTZ COWAN

As the Civil War came to an end, America began the transition from a preindustrial society into the industrial age. The period from 1870 to 1920 saw many changes in America, and at the end of that era, it had become the largest economy in the world—one that was far less dependent on nature than it once had been.

While we currently use many different technological systems, Ruth Schwartz Cowan, professor emeritus of history at the State University of New York, Stony Brook, discusses five developed during that fifty-year period to demonstrate the increasing complexity of life and the development of infrastructure we now largely take for granted. Eventually, this increasing complexity spread across the world. The characteristics of an industrialized society can now be found even in the most remote corners of the globe.

The increased levels of productivity in manufacturing, including agriculture, had economic as well as political ramifications. As these changes occurred, few thought about what was happening or what the long-range implications would be. Today's globalized world is vastly different from what it once was, with many people more reliant and interdependent on one another. Our lives are increasingly more intertwined than ever before.

FOCUS QUESTIONS

1. Explain the rationale behind Cowan's assertion that industrialization has made us less independent and more closely connected to many other people than ever before.
2. Consider the key points of this article in relation to the ethical issues raised in the Jonas piece (Selection 1.3.1). What new individual responsibilities might arise from the new networks extending far beyond the nuclear family?
3. If technology continues to improve the quality of life for many more people in the world, what might our future look like if the networks discussed by Cowan become larger and more common, linking even greater numbers of people throughout the world?

KEYWORDS

balance of trade, entrepreneurs, hunter-gatherers, industrialization, infrastructure, international trade, investment banking, manufacturing, productivity, transportation

Source: A SOCIAL HISTORY OF AMERICAN TECHNOLOGY by Cowan (1997) pp. 149–172
© 1997 by Oxford University Press, Inc. By permission of Oxford University Press, USA.

Between 1870 and 1920, the United States changed in ways that its founders could never have dreamed possible. Although American industrialization began in the 1780s, the nation did not become an industrialized society until after the Civil War had ended. The armistice agreed to at Appomattox signaled, although the participants probably did not realize it, the beginning of the take-off phase of American industrialization. Having begun as a nation of farmers, the United States became a nation of industrial workers. Having begun as a financial weakling among the nations, by 1920 the United States had become the world's largest industrial economy.

What did this transformation mean to the people who lived through it? When a society passes from preindustrial to industrial conditions, which is what happened in the United States in the years between 1870 and 1920, people become less dependent on nature and more dependent on each other. This is one of history's little ironies. In a preindustrial society, when life is unstable, the whims of the weather and the perils of natural cycles are most often to blame. In an industrial society, when life is unstable, the whims of the market and the perils of social forces are most often to blame. Put another way, this means that in the process of industrialization individuals become more dependent on one another because they are linked together in large, complex networks that are, at one and the same time, both physical and social: technological systems.

INDUSTRIALIZATION, DEPENDENCY, AND TECHNOLOGICAL SYSTEMS

Many Americans learned what it means to become embedded in a set of technological systems in the years between 1870 and 1920. Today we have become so accustomed to these systems that we hardly ever stop to think about them; although they sustain our lives, they nonetheless remain mysterious. In the late twentieth century, people have tended to think that, if anything, industrialization has liberated them from dependency, not encased them in it, but that is not the case. We can see this clearly by imagining how a woman might provide food for a two-year-old child in a nonindustrialized society.

In a hunter-gatherer economy, she might simply go into the woods and collect nuts or walk to the waterside and dig for shellfish. In a premodern agricultural community (such as the one that some of the native peoples of the eastern seaboard had created), she might work with a small group of other people to plant corn, tend it, harvest it, and shuck it. Then she herself might dry it, grind it into meal, mix it with water, and bake it into bread for the child to eat. In such a community, a woman would be dependent on the cooperation of several other people in order to provide enough food for her child, but all of those people would be known to her and none of them would be involved in an activity in which she could not have participated if necessity had demanded.

In an industrialized economy (our own, for example), an average woman's situation is wholly different. In order to get bread for a child, an average American woman is dependent on thousands of other people, virtually all of them totally unknown to her, many of them living and working at a considerable distance, employing equipment that she could not begin to operate, even if her life (quite literally) depended on it and even if she had the money (which isn't likely) to purchase it. A farmer grew the wheat using internal combustion engines and petroleum-derivative fertilizers. Then the wheat was harvested and transported to an organization that stored it under stable conditions, perhaps for several years. Then a milling company may have purchased it and transported it (over thousands of miles of roads or even ocean) to a mill, where it was ground by huge rollers powered by electricity (which itself may have been generated thousands of miles away). Then more transportation (all of this transportation required petroleum, which itself had to be processed and transported) was required: to a baking factory, where dozens of people (and millions of dollars of machinery)

were used to turn the flour into bread. Then transportation again: to a market, where the woman could purchase it (having gotten herself there in an automobile, which itself had to be manufactured somewhere else, purchased at considerable expense, and supplied with fuel)—all of this before a slice of it could be spread with peanut butter to the delight of a two-year-old.

The point should, by now, be clear. People who live in agricultural societies are dependent on natural processes: they worry, with good reason, about whether and when there will be a drought or a flood, a plague of insects or of fungi, good weather or bad. People who live in industrial societies are not completely independent of such natural processes, but are more so than their predecessors (many floodplains have been controlled; some droughts can be offset by irrigation). At the same time, they are much more dependent on other people and on the technological systems that other people have designed and constructed. The physical parts of these systems are networks of connected objects: tractors, freight cars, pipelines, automobiles, display cases. The social parts are networks of people and organizations that make the connections between objects possible: farmers, bakers, and truck drivers; grain elevators, refineries, and supermarkets.

Preindustrialized societies had such networks of course, but in industrialized societies, the networks are more complex and much denser—all of which makes it much harder for individuals to extricate themselves. A small change very far away can have enormous effects very quickly. Daily life can be easily disrupted for reasons that ordinary people can find hard to understand, and even experts can have difficulty comprehending.

People live longer and at a higher standard of living in industrial societies than in preindustrial ones, but they are not thereby rendered more independent (although advertising writers and politicians would like them to think they are) because, in the process of industrialization, one kind of dependency is traded for another: nature for technology. Americans learned what it meant to make that trade in the years between 1870 and 1920. We can begin understanding what they experienced if we look at some of the technological systems that were created or enlarged during those years.

THE TELEGRAPH SYSTEM

The very first network that Americans experienced really looked like a network: the elongated spider's web of electric wires that carried telegraph signals. The fact that electricity could be transmitted long distances through wires had been discovered in the middle of the eighteenth century. Once a simple way to generate electric currents had been developed (a battery, or voltaic pile, named after the man who invented it, Alessandro Volta), many people began experimenting with various ways to send messages along the wires. An American portrait painter, Samuel F. B. Morse, came up with a practicable solution. Morse developed a transmitter that emitted a burst of electric current of either short or long duration (dots and dashes). His receiver, at the other end of the wire, was an electromagnet, which, when it moved, pushed a pencil against a moving paper tape (thus recording the pattern of dots and dashes). The most creative aspect of Morse's invention was his code, which enabled trained operators to make sense out of the patterns of dots and dashes.

In 1843, after Morse had obtained a government subvention, he and his partners built the nation's first telegraph line between Baltimore and Washington. By 1845, Morse had organized his own company to build additional lines and to license other telegraph companies so that they could build even more lines, using the instruments he had patented. In a very short time, however, dozens of competing companies had entered the telegraph business, and Morse had all he could do to try to collect the licensing fees to which he was entitled. By 1849, almost every state east of the Mississippi had telegraph service, much of it provided by companies that were exploiting Morse's patents without compensating him.

Beginning around 1850, one of these companies, the New York and Mississippi Valley Printing Telegraph Company, began buying up or merging with all the others; in 1866, it changed its name to

the Western Union Telegraph Company. In the decades after the Civil War, Western Union had an almost complete monopoly on telegraph service in the United States; a message brought to one of its offices could be transmitted to any of its other offices in almost all fairly large communities in the United States. Once the message was delivered, recipients could pick it up at a Western Union office. During these decades, only one company of any note succeeded in challenging Western Union's almost complete monopoly on telegraph service. The Postal Telegraph Company specialized in providing pick-up and delivery services for telegrams; yet even at the height of its success, it never managed to corner more than 25 percent of the country's telegraph business.

In 1866, when Western Union was incorporated, it already controlled almost 22,000 telegraph offices around the country. These were connected by 827,000 miles of wire (all of it strung from a virtual forest of telegraph poles, many of them running along railroad rights of way), and its operators were handling something on the order of 58 million messages annually. By 1920, the two companies (Western and Postal) between them were managing more than a million miles of wire and 155 million messages. Yet other companies (many of the railroads, for example, several investment banking houses, several wire news services) were using Western Union and Postal Telegraph lines on a contractual basis to provide in-house communication services (the famous Wall Street stock ticker was one of them).

As a result, as early as 1860, and certainly by 1880, the telegraph had become crucial to the political and economic life of the nation. Newspapers had become dependent on the telegraph for quick transmission of important information. The 1847 war with Mexico was the first war to have rapid news coverage, and the Civil War was the first in which military strategy depended on the quick flow of battle information over telegraph lines. During the Gilded Age (1880–1900), the nation's burgeoning financial markets were dependent on the telegraph for quick transmission of prices and orders. Railroad companies used the telegraph for scheduling and signaling purposes since information about deviations in train times could be quickly transmitted along the lines. The central offices of the railroads utilized telegraph communication to control the financial affairs of their widely dispersed branches. When the Atlantic cable was completed in 1866, the speed and frequency of communication between nations increased, thereby permanently changing the character of diplomatic negotiations. The cable also laid the groundwork for the growth of international trade (particularly the growth of multinational corporations) in the later decades of the century.

In short, by 1880, if by some weird accident all the batteries that generated electricity for telegraph lines had suddenly run out, the economic and social life of the nation would have faltered. Trains would have stopped running; businesses with branch offices would have stopped functioning; newspapers could not have covered distant events; the president could not have communicated with his European ambassadors; the stock market would have had to close; family members separated by long distances could not have relayed important news—births, deaths, illnesses—to each other. By the turn of the century, the telegraph system was both literally and figuratively a network, linking together various aspects of national life—making people increasingly dependent on it and on one another.

THE RAILROAD SYSTEM

Another system that linked geographic regions, diverse businesses, and millions of individuals was the railroad. We have already learned about the technical developments (the high-pressure steam engine, the swivel truck, the T-rail) that were crucial to the development of the first operating rail lines in the United States in the 1830s. Once the technical feasibility of the railroad became obvious, its commercial potential also became clear. The railroad, unlike canals and steamboats, was not dependent on proximity to waterways and was not (as boats were) disabled when rivers flooded or canals froze.

During the 1840s, American entrepreneurs had began to realize the financial benefits that

railroading might produce and railroad-building schemes were being concocted in parlors and banks, state houses, and farm houses all across the country. By the 1850s, a good many of those schemes had come to fruition. With 9,000 miles of railroad track in operation, the United States had more railroad mileage than all other western nations combined; by 1860, mileage had more than trebled, to 30,000 miles.

The pre-Civil War railroad system was not yet quite a technological system because, large as it was, it still was not integrated as a network. Most of the existing roads were short-haul lines, connecting such major cities as New York, Chicago, and Baltimore with their immediate hinterlands. Each road was owned by a different company, each company owned its own cars, and each built its tracks at the gauge (width) that seemed best for the cars it was going to attempt to run and the terrain over which the running had to be done. This lack of integration created numerous delays and additional expenses. In 1849, it took nine transshipments between nine unconnected railroads (and nine weeks of travel) to get freight from Philadelphia to Chicago. In 1861, the trip between Charleston and Philadelphia required eight car changes because of different gauges. During and immediately after the Civil War, not a single rail line entering either Philadelphia or Richmond made a direct connection with any other, much to the delight of the local teamsters, porters, and tavern keepers.

The multifaceted processes summed up under the word "integration" began in the years just after the Civil War and accelerated in the decades that followed. The rail system grew ever larger, stretching from coast to coast (with the completion of the Union Pacific Railroad in 1869), penetrating into parts of the country where settlement did not yet even exist. There were roughly 53,000 miles of track in 1870, but there were 93,000 miles by the time the next decade turned, and 254,000—the all-time high—by 1920. In that half century, the nation's population tripled, but its rail system grew sevenfold; the forty-eight states of the mainland United States became phys-

The form of the rail system was just as significant as its size. By 1920, what had once been a disjointed collection of short (usually north–south) lines had been transformed into a network of much longer trunk lines (running from coast to coast, east-west), each served by a network of shorter roads that connected localities (the limbs) with the trunks. Passengers could now travel from New York to San Francisco with only an occasional change of train and freight traveled without the necessity of transshipments. What had made this kind of integration possible was not a technological change, but a change in the pattern of railroad ownership and management.

From the very beginning of railroading, railroad companies had been joint-stock ventures. Huge amounts of capital had been required to build a railroad: rights of way had to be purchased, land cleared, bridges built, locomotives ordered, passenger cars constructed, freight cars bought. Once built, railroads were very expensive to run and to maintain: engines had to be repaired, passengers serviced, freight loaded, tickets sold, stations cleaned. Such a venture could not be financed by individuals, or even by partnerships. Money had to be raised both by selling shares of ownership in the company to large numbers of people and by borrowing large sums of money by issuing bonds.

As a result, both American stockbroking and American investment banking were twin products of the railroad age. Some of America's largest nineteenth-century fortunes were made by people who knew not how to build railroads, but how to finance them: J. P. Morgan, Leland Stanford, Jay Gould, Cornelius Vanderbilt, and George Crocker. These businessmen consolidated the railroads. They bought up competing feeder lines; they sought control of the boards of directors of trunk lines; they invested heavily in the stock of feeder roads until the feeders were forced to merge with the trunks. When they were finished, the railroads had become an integrated network, a technological system. In 1870, there had been several hundred railroads, many of which were in direct competition with each other. By 1900, virtually all the railroad mileage in the United States

(often mutually cooperative) railroad combinations, all of which owed their existence to the machinations of a few very wealthy investment bankers.

As railroad ownership became consolidated, the railroad system became physically integrated. The most obvious indicator of this integration was the adoption of a standard gauge, which made it unnecessary to run different cars on different sets of tracks. By the end of the 1880s, virtually every railroad in the country had voluntarily converted to a gauge of 4 feet, 8½ inches in order to minimize both the expense and the delays of long-distance travel. On this new integrated system, the need for freight and passengers to make repeated transfers was eliminated; as a result, costs fell while transportation speed increased.

The railroad system had a profound impact on the way in which Americans lived. By 1900, the sound of the train whistle could be heard in almost every corner of the land. Virtually everything Americans needed to maintain and sustain their lives was being transported by train. As much as they may have grumbled about freight rates on the railroads (and there was much injustice, particularly to farmers, to grumble about) and as much as they may have abhorred the techniques that the railroad barons had used to achieve integration, most Americans benefited from the increased operational efficiency that resulted.

In the years in which population tripled and rail mileage increased seven times, freight tonnage on the railroads went up elevenfold. Cattle were going by train from the ranches of Texas to the slaughter houses of Chicago; butchered beef was leaving Chicago in refrigerated railroad cars destined for urban and suburban kitchens. Lumber traveled from forests to sawmills by train; two-by-four beams to build houses on the treeless plains left the sawmills of the Pacific Northwest on flat-cars. Some petroleum went from the well to the refinery by train; most kerosene and gasoline went from the refinery to the retailer by train. Virtually all the country's mail traveled by train, including cotton cloth and saddles, frying pans and furniture ordered from the mail-order compa-

Even as fundamental and apparently untransportable a commodity as time was affected by the integration of the rail system, for scheduling was an important facet of integration. People who were going to travel by train had to know what time their trains would leave, and if connections had to be made, trains had to be scheduled so as to make the connections possible. Schedules also had to be constructed, especially on heavily trafficked lines, to ensure that trains did not collide. But scheduling was exceedingly difficult across the long distances of the United States because communities each established their own time on the basis of the position of the sun. When it was noon in Chicago, it was 12:30 in Pittsburgh (which is to the east of Chicago) and 11:30 in Omaha (to the west). The train schedules printed in Pittsburgh in the early 1880s listed six different times for the arrival and departure of each train. The station in Buffalo had three different clocks.

Sometime in the early 1880s, some professional railroad managers and the editors of several railroad publications agreed to the idea, first proposed by some astronomers, that the nation should be divided into four uniform time zones. By common agreement among the managers of the country's railroads, at noon (in New York) on Sunday, November 18, 1883, railroad signalmen across the country reset their watches. The zones were demarcated by the 75th, 90th, 105th, and 120th meridians. People living in the eastern sections of each zone experienced, on that otherwise uneventful Sunday, two noons, and people living in the western sections skipped time. Virtually everyone in the country accepted the new time that had been established by the railroads, although Congress did not actually confirm the arrangement by legislation for another thirty-five years. Such was the pervasive impact of the integrated rail network.

THE PETROLEUM SYSTEM

In 1859, a group of prospectors dug a well in a farmyard in Titusville, Pennsylvania. Although they appeared to be looking for water, the prospectors were in fact searching for an under-

had been bubbling to the surface of nearby land and streams. Native Americans had used this combustible substance as a lubricant for centuries. The prospectors were hoping that if they could find a way to tap into an underground reservoir of this material, they could go into the business of selling it to machine shops and factories (as a machine lubricant, an alternative to animal fat) and to households and businesses (as an illuminant, an alternative to whale oil and candles).

The prospectors struck oil—and the American petroleum industry was born. Within weeks the news had spread, and hundreds of eager profiteers rushed into western Pennsylvania, hoping to purchase land, drill for oil, or find work around the wells. The Pennsylvania oil rush was as massive a phenomenon as the California gold rush a decade earlier.

The drillers soon discovered that crude petroleum is a mixture of oils of varying weights and characteristics. These oils, they learned, could be easily separated from one another by distillation, an ancient and fairly well-known craft. All that was needed was a fairly large closed vat with a long outlet tube (called a still) and a fire. The oil was heated in the still and the volatile gases produced would condense in the outlet tube. A clever distiller (later called a refiner) could distinguish different portions (fractions) of the distillate from each other, and then only the economically useful ones needed to be bottled and sent to market.

The market for petroleum products boomed during the Civil War: northern factories were expanding to meet government contracts; the whaling industry was seriously hampered by naval operations; railroads were working overtime to transport men and material to battlefronts. By 1862, some 3 million barrels of crude oil were being processed every year. Under peacetime conditions the industry continued to expand; by 1872, the number of processed barrels had trebled.

Transportation of petroleum remained a problem, however. The wells were located in the rural, underpopulated Appalachian highlands of Pennsylvania, not only many miles away from the cities in which the ultimate consumers lived, but also many miles away from railroad lines that served those cities. Initially crude oil had been collected in barrels and had been moved (by horse and cart or by river barges) to railroad-loading points. There the barrels were loaded into freight cars for the trip to the cities (such as Cleveland and Pittsburgh) in which the crude was being refined and sold. The transportation process was cumbersome, time-consuming, and wasteful; the barrels leaked, the barges sometimes capsized, the wagons—operating on dirt roads—sometimes sank to their axles in mud.

Pipelines were an obvious solution, but a difficult one to put into practice, given that no one had ever before contemplated building and then maintaining a continuous pipeline over the mountainous terrain and the long distances that had to be traversed. The first pipeline to operate successfully was built in 1865. Made of lap-welded cast-iron pipes, two inches in diameter, it ran for six miles from an oil field to a railroad loading point and had three pumping stations along the way. This first pipeline carried eighty barrels of oil an hour and had demonstrated its economic benefits within a year. Pipeline mileage continued to increase during the 1870s and 1880s (putting thousands of teamsters out of business), but virtually all of the lines were relatively short hauls, taking oil from the fields to the railroads. Throughout the nineteenth century and well into the twentieth, the railroads were still the principal long-distance transporters of both crude and refined oil. After the 1870s, the drillers, refiners, and railroads gradually dispensed with barrels (thus putting thousands of coopers out of business) and replaced them with specially built tank cars, which could be emptied into and loaded from specially built holding tanks. As it was being constructed, the network of petroleum pipelines was thus integrated into the network of railroad lines. It was also integrated into the telegraph network. Oil refineries used the telegraph system partly to keep tabs on prices for oil in various localities and partly to report on the flow of oil through the lines.

The most successful petroleum entrepreneurs were the ones who realized that control of petroleum transportation was the key ingredient in control of the entire industry. The major actor

in this particular economic drama was John D. Rockefeller. Rockefeller had been born in upstate New York, the son of a talented patent medicine salesman, but he had grown up in Cleveland, Ohio, a growing commercial center (it was a Great Lake port and both a canal and railroad terminus), and had learned accountancy in a local commercial college. His first job was as a bookkeeper for what was then called a commission agent, a business that collected commissions for arranging the shipment of bulk orders of farm products. A commission agent's success depended on getting preferential treatment from railroads and shipping companies. Rockefeller carried this insight with him, first when he went into a partnership as his own commission agent and then, in 1865, when he became the co-owner of an oil refinery in Cleveland.

Rockefeller and his associates were determined to control the then chaotic business of oil refining. They began by arranging for a secret rebate of oil shipments from one of the two railroads then serving Cleveland. Then in the space of less than a month, using the rebate as an incentive, they managed to coerce other Cleveland refiners into selling out and obtained control of the city's refining. Within a year or two, Rockefeller was buying up refineries in other cities as well. He had also convinced the railroads that he was using that they should stop carrying oil to refineries owned by others, so that he was in almost complete control of the price offered to drillers. In the early 1870s, a group of drillers banded together to build pipelines that would take their oil to railroads with which Rockefeller wasn't allied. Rockefeller responded to this challenge by assembling a monopoly on the ownership of tank cars (since the pipelines did not go all the way to the refineries and railroad tank cars were still necessary), and by 1879, he had been so successful in squeezing the finances of the pipeline companies that their stockholders were forced to sell out to him. In that year, as a result of their control both of refineries and pipelines, Rockefeller and his associates controlled 90 percent of the refined oil in the United States.

Having bought up the competing pipelines (having let other people take the risks involved in developing new technologies for building and maintaining those lines), Rockefeller was quick to see their economic value. In 1881, one of his companies completed a six-inch line from the Pennsylvania oil fields to his refinery in Bayonne, New Jersey—the first pipeline that functioned independently of the railroads. By 1900, Rockefeller had built pipelines to Cleveland, Philadelphia, and Baltimore, and Standard Oil (Rockefeller's firm) was moving 24,000 barrels of crude a day (he still used the railroads to move the oil after it had been refined).

By that point, hundreds of civil and mechanical engineers were working for Rockefeller's pipeline companies (which held several patents on pipeline improvements), and several dozen chemists and chemical engineers were working in his refineries (and developing new techniques, such as the Frasch process for taking excess sulfur out of petroleum). In addition, Standard Oil was pioneering financial, management, and legal techniques for operating a business that had to control a huge physical network, spread out over several states. Since the laws dealing with corporations differed in each state and since some of them prevented a corporation in one state from owning property in another, one of Rockefeller's attorneys worked out a corporate arrangement so that Standard Oil had a different corporation in each state in which it operated (Standard Oil of New Jersey, Standard Oil of Ohio, and so forth). The stockholders in each corporation turned their stock over to a group of trustees, who managed the whole enterprise from New York—the famous Standard Oil Trust, of which Rockefeller himself was the single largest stockholder and therefore the major trustee. (The trust, as a way to organize a complex business, was soon picked up in tobacco and sugar refining and other industries involved in large-scale chemical processing, leading Congress, worried about the monopolistic possibilities, to pass the Sherman Anti-Trust Act in 1890.)

By 1900, the Standard Oil Trust (which had successfully battled antitrust proceedings in court) controlled most of the oil produced in Pennsylvania, and it owned most of the new oil fields

Rockefeller's almost complete stranglehold on the industry wasn't broken until oil was discovered early in the twentieth century in Texas, Oklahoma, Louisiana, and California, outside the reach of the pipelines he controlled and the railroads with which he was associated. Increased competition was accompanied by the continued growth not only of the pipeline network, but also of the industry as a whole: 26 million barrels of petroleum were processed in 1880, 45 million in 1890, 63 million in 1900, 209 million in 1910 (as gasoline was just beginning to edge out kerosene as the most important petroleum product), and 442 million in 1920 (when the Model T had been in production for almost eight years).

Like the telegraph and the railroad (and in combination with the telegraph and the railroad), the oil pipeline network had become a pervasive influence on the American economy and on the daily life of Americans. In the last decades of the nineteenth century, a very large number of Americans, especially those living outside of the major cities, used one of its products, kerosene, for heating and lighting their homes and for cooking. During the same decades, American industry became dependent on other fractions of petroleum to lubricate the machinery with which it was producing everything from luxurious cloth to common nails. Finally, in the early decades of the twentieth century, with the advent first of the internal combustion engine fueled by gasoline and then of automobiles and trucks powered by that engine, Americans discovered that access to petroleum was becoming a necessary condition not only of their working lives but also of their leisure time.

THE TELEPHONE SYSTEM

Technologically, the telephone was similar to the telegraph, but socially it was very different. The device patented by Alexander Graham Bell in 1876 was rather like a telegraph line: voices rather than signals could be transmitted by electric current because the transmitter lever and the receiving pencil had been replaced by very sensitive diaphragms. Aware of the difficulties that Morse had encountered in reaping profits from his patents—and aware that he had no head for business—Bell decided to turn over the financial and administrative details of creating a telephone network to someone else.

The businessmen and the attorneys who managed the Bell Telephone Company did their work well. While the railroad, telegraph, and petroleum networks had been integrated by corporate takeovers, the telephone system was integrated, from the very beginning, by corporate design. A crucial decision had been made early on: Bell Telephone would manufacture all the telephone instruments, then lease the instruments to local companies, which would operate telephone exchanges under license to Bell. This meant that for the first sixteen years of telephone network development (sixteen years was then the length of monopoly rights under a patent), the Bell Telephone Company could dictate, under the licensing agreements, common technologies for all the local telephone systems. Bell could also control the costs of telephone services to local consumers.

Because of this close supervision by one company, the telephone system was integrated from the very beginning. Between 1877 and 1893, the Bell Telephone Company, through its affiliated local operating companies, controlled and standardized virtually every telephone, every telephone line, and every telephone exchange in the nation. Indeed in the 1880s, the officers of Bell were confident that they could profitably begin long-distance service (that is, service that would connect one local operating company with another) precisely because all of the operating companies were using its standardized technology. Bell needed to hire physicists and electrical engineers to solve the technical problems involved in maintaining voice clarity over very long wires, but the organizational problems involved in connecting New York with Chicago and Chicago with Cleveland turned out to be minimal.

On the assumption that the telephone system would end up being used very similarly to the telegraph network, the officers of Bell had decided that their most important customers would be other businesses, particularly those in urban areas. They decided, as a marketing strategy, to keep rates fairly high, in return for which they

reliable service possible. By the end of the company's first year of operation, 3,000 telephones had been leased, 1 for every 10,000 people. By 1880, there were 60,000 (1 per 1,000), and when the Bell patents expired in 1893, there were 260,000 (1 per 250). About two-thirds of these phones were located in businesses. Most of the country's business information was still traveling by mail and by telegraph (because businessmen wanted a written record of their transactions), but certain kinds of businesses were starting to find the telephone very handy: in 1891, the New York and New Jersey Telephone Company served 937 physicians and hospitals, 401 pharmacies, 363 liquor stores, 315 stables, 162 metalworking plants, 146 lawyers, 126 contractors, and 100 printing shops.

After the Bell patents expired, independent telephone companies entered the business despite Bell's concerted effort to keep them out. By 1902, there were almost 9,000 such independent companies, companies not part of the Bell system. When the organizers of the Bell system had analogized the telephone to the telegraph, they had made a crucial sociological mistake. They understood that in technological terms the telephone was similar to the telegraph, but they failed to understand that in social terms it was quite different. The telephone provided user-to-user communication (with the telegraph there were always intermediaries). In addition, the telephone was a form of voice communication; it facilitated emotional communication, something that was impossible with a telegraph. In short, what the organizers of the Bell system had failed to understand was that people would use the telephone to socialize with each other.

The independent companies took advantage of Bell's mistake. Some of them offered services that Bell hadn't thought to provide. Dial telephones were one such service, allowing customers to contact each other without having to rely on an operator (who sat at a switchboard, manually connecting telephone lines, one to another, with plugs). Operators were notorious for relieving the boredom of their jobs by listening in on conversations, something many customers wanted to avoid. Party lines were another such service. Anywhere from two to ten residences could share the same telephone line and telephone number, which drastically lowered the costs of residential services. Many lower-income people turned out to be willing to put up with the inconvenience of having to endure the ringing of telephones on calls meant for other parties in exchange for having telephone service at affordable rates.

Yet other independent companies served geographic locales that the Bell companies had ignored. This was particularly the case in rural areas where there were farm households. Bell managers apparently hadn't thought that farmers would want telephones, but it turned out that they were wrong. Farm managers used telephones to get prompt reports on prices and weather. Farm households used telephones to summon doctors in emergencies and to alleviate the loneliness of lives lived far from neighbors and relatives. In 1902, relatively few farm households had telephones, but as the independent companies grew, so did the number of farm-based customers; by 1920, just under 39 percent of all farm households in the United States had telephone service (while only 34 percent of non-farm households did).

All this competition in telephone service had the net effect that any economist could have predicted: prices for telephone service fell, even in the Bell system. In order to keep the system companies competitive, the central Bell company had to cut the rates that it charged its affiliates for the rental of phones, and these savings were passed on to consumers. In New York City, as just one example, rates fell from $150 for 1,000 calls in 1880 to $51 in 1915 (figures adjusted for inflation).

As a result, in the period between 1894 and 1920, the telephone network expanded profoundly. Middle-class people began to pay for telephone service to their homes. Farm households became part of the telephone network (in record numbers). Retail businesses began to rely on telephones in their relations with their customers. By 1920, there were 13 million telephones in use in the country, 123 for every 1,000 people. Eight million of those 13 million phones belonged to Bell and 4 million to independent companies

that connected to Bell lines. In just forty years, the telephone network, which provided point-to-point voice communication, had joined the telegraph, railroad, and petroleum networks as part of the economic and social foundation of industrial society.

THE ELECTRIC SYSTEM

Like the telegraph and telephone systems, the electric system was (and still is) quite literally a network of wires. Physicists, who had been experimenting with electricity since the middle of the eighteenth century, knew that under certain conditions electricity could produce light. Unfortunately, the first devices invented for generating a continuous flow of electricity—batteries—did not create a current strong enough for illumination. However, in 1831 the British experimenter Michael Faraday perfected a device that was based on a set of observations that scientists had made a decade earlier: all electric current will make a magnet move and a moving magnet will create an electric current. Faraday built an electric generator (a rotating magnet with a conducting wire wound around it)—a device that could, unlike the battery, create a continuous flow of current strong enough to be used for lighting.

Within a short time, the generator was being used to power arc lamps in which the light (and a lot of heat) was produced by sparking across a gap in the conducting wires. Arc lamps were first used in British and French lighthouses in the 1860s; the generator that created the electricity was powered by a steam engine. A few years later, arc lamps were also being used for street lighting in some American cities. Unfortunately, arc lamps were dangerous; they had to be placed very far away from people and from anything that might be ignited by the sparks. By the mid-1870s, several people in several different countries were racing with each other to find a safer form of electrical lighting, the incandescent lamp. In such a lamp, light would be derived from a glowing, highly resistant filament and not a spark; but the filament had to be kept in a vacuum so that it wouldn't oxidize (and disappear) too fast.

Thomas Alva Edison won the race. In 1878, he already had amassed a considerable reputation (and a moderate fortune) as an inventor. His first profitable invention had been the quadruplex telegraph, which could carry four messages at once, and he had also made successful modifications to the stock ticker, the telegraph system for relaying stock prices from the floor of the stock exchange to the offices of investors and brokers. These inventions had enhanced his reputation with Wall Street financiers and attorneys. In 1876, when he decided to become an independent inventor, building and staffing his own laboratory in Menlo Park, New Jersey, and again in 1878, when he decided that he wanted his laboratory to crack the riddle of electric lighting, he had no trouble borrowing money to invest in the enterprise.

Actually, they were enterprises. From the beginning, Edison understood that he wanted to build a technological system *and* a series of businesses to manage that system. The first of these businesses was the Edison Electric Light Company, incorporated for the purpose of financing research and development of electric lighting. Most of the stock was purchased by a group of New York financiers; Edison received stock in return for the rights to whatever lighting patents he might develop. Once Edison had actually invented a workable lightbulb (it had a carbonized thread as its filament), he proceeded to design other devices, and create other companies, that would all be parts of the system. The Edison Electric Illuminating Company of New York, founded in 1880, was created to build and maintain the very first central generating station providing electric service to customers. When this station opened its doors in 1882 (as its site Edison chose the part of Manhattan with the highest concentration of office buildings), it contained several steam-driven generators (built to Edison's design by the Edison Machine Company) and special cables to carry the electricity underground (made by the Edison Electric Tube Company). Customers who signed up for electric service had their usage measured by meters that Edison had invented; their offices were outfitted with lamp sockets that Edison had designed into which they were to place lightbulbs that another Edison

Information about this new system spread very fast (thanks to publicity generated by the Edison Electric Light Company), and within a few months (not even years), entrepreneurs were applying to Edison for licenses to build electric generating plants all over the country, indeed all over the world. Having been designed as a system, the electrical network grew very fast. There was only one generating plant in the country in 1882, but by 1902, there were 2,250, and by 1920, almost 4,000. These plants had a total generating capacity of 19 million kilowatts. Just over a third of the nation's homes were wired for electricity by 1920, by which time electricity was being used not only for lighting but also for cooling (electric fans), ironing (the electric iron replaced the so-called sad iron quickly), and vacuuming (the vacuum cleaner was being mass-produced by 1915).

The Edison companies (some of which eventually merged with other companies to become the General Electric Company) were not, however, able to remain in control of the electric system for as long (or as completely) as the Bell companies were able to dominate the telephone business or Standard Oil the petroleum business. Part of the reason for this lay in the principles of electromagnetic induction, which can be used to create electric motors as well as electric generators. The same experimenters who were developing electric generators in the middle years of the nineteenth century were also developing electric motors, and one of the first applications of those motors was in a business very different from the lighting business: electric traction for electric intraurban streetcars, often known as trolley cars. The first of these transportation systems was installed in Richmond, Virginia, in 1888 by a company owned by Frank Sprague, an electrical engineer who had briefly worked for Edison.

Sprague had invented an electric motor that, he thought, would be rugged enough to power carriages running day in and day out on city streets. As it turned out, the motor had to be redesigned, and redesigned again, before it worked very well, and Sprague also had to design trolley poles (for conducting the electricity from

controlling system (so that the speed of the motor could be varied by the person driving the carriage). In the end, however, the electric streetcar was successful, and the days of the horse-pulled carriage were clearly numbered. Fourteen years after Sprague's first system began operating, the nation had 22,576 miles of track devoted to street railways.

Electric motors were also being used in industry. The earliest motors, like the streetcar motors, had been direct current (D.C.) motors, which needed a special and often fragile device (called a commutator) to transform the alternating current (A.C.) produced by generators. In 1888, an A.C. motor was invented by Nikola Tesla, a Serbian physicist who had emigrated to the United States. Tesla's patents were assigned to the Westinghouse Company, which began both to manufacture and to market them. At that point, the use of electric motors in industry accelerated. The very first factory to be completely electrified was a cotton mill, built in 1894. As electric motors replaced steam engines, factory design and location changed; it was no longer necessary to build factories that were several stories high (to facilitate power transmission from a central engine) or to locate them near water sources (to feed the steam boilers). The first decade of the twentieth century was a turning point in the use of electric power in industry as more and more factories converted; by 1901, almost 400,000 motors had been installed in factories, with a total capacity of almost 5 million horsepower.

In short, the electrical system was more complex than the telephone and petroleum systems because it consisted of several different subsystems (lighting, traction, industrial power) with very different social goals and economic strategies; because of its complexity, no single company could dominate it. By 1895, when the first generating plant intended to transmit electricity over a long distance became operational (it was a hydroelectric plant built to take advantage of Niagara Falls, transmitting electricity twenty miles to the city of Buffalo), there were several hundred companies involved in the electric industry: enormous companies such as Westinghouse and General

to lightbulbs; medium-sized companies, such as the ones that ran streetcar systems or that provided electric service to relatively small geographic areas; and small companies, which made specialized electric motors or parts for electric motors. Despite this diversity, the electric system was unified by the fact that its product, electric energy, had been standardized. By 1910, virtually all the generating companies (which, by now, had come to be called utility companies) were generating alternating current at sixty cycles per second. This meant that all electric appliances were made to uniform specifications and all transmission facilities could potentially be connected to one another. By 1920, electricity had supplanted gas, kerosene, and oils for lighting. In addition, it was being used to power sewing machines in ready-made clothing factories, to separate aluminum from the contaminants in its ores, to run projectors through which motion pictures could be viewed, to carry many thousands of commuters back and forth, and to do dozens of other chores in workplaces and residences. As transmission towers marched across the countryside and yet another set of wire-carrying poles were constructed on every city street, few Americans demonstrated any inclination to decline the conveniences that the youngest technical system—electricity—was carrying in its wake.

THE CHARACTER OF INDUSTRIALIZED SOCIETY

As inventors, entrepreneurs, and engineers were building all these multifarious technological systems, Americans were becoming increasingly dependent on them. Each time a person made a choice—to buy a kerosene lamp or continue to use candles, to take a job in an electric lamp factory or continue to be a farmer, to send a telegraph message instead of relying on the mail, to put a telephone in a shop so that customers could order without visiting—that person, whether knowingly or not, was becoming increasingly enmeshed in a technological system. The net effect of all that construction activity and all those choices was that a wholly new social order, and wholly different set of social and economic relationships between people, emerged: industrial society.

In industrial societies, manufactured products play a more important economic role than agricultural products. More money is invested in factories than in farms; more bolts of cloth are produced than bales of hay; more people work on assembly lines than as farm laborers. Just over half (53 percent) of what was produced in the United States was agricultural in 1869 and only a third (33 percent) was manufactured. In 1899 (just thirty years later), those figures were reversed: half the nation's output was in manufactured goods and only a third was agricultural, despite the fact that the nation's total farm acreage had increased rapidly as a result of westward migration. Manufacturing facilities were turning out products that were becoming increasingly important aspects of everyday life: canned corn and lightbulbs, cigarettes and underwear.

In a preindustrial society, the countryside is the base for economic and political power. In such societies, most people live in rural districts. Most goods that are traded are agricultural products; the price of fertile land is relatively high; and wealth is accumulated by those who are able to control that land. Industrialized societies are dominated by their cities. More people live and work in cities than on farms; most goods are manufactured in cities; most trade is accomplished there; wealth is measured in money and not in land. Furthermore, the institutions that control money—banks—are urban institutions.

As the nineteenth century progressed, more and more Americans began living either in the rural towns in which factories were located (which, as a result, started to become small cities) or in the older cities that had traditionally been the center of artisanal production and of commerce. Native-born Americans began moving from the countryside to the city; many newly arrived Americans (and there were millions of newcomers to America in the nineteenth century) settled in cities. Just over half of all Americans (54 percent) were farmers or farm laborers in 1870, but only one in three was by 1910. Some American families underwent the rural-urban transition slowly: a daughter might move off the farm

to a rural town when she married, and then a granddaughter might make her fortune in a big city. Others had less time: a man might be tending olive groves in Italy one day and working in a shoe factory in Philadelphia two months later.

During the 1840s, the population of the eastern cities nearly doubled, and several mid-western cities (St. Louis, Chicago, Pittsburgh, Cincinnati) began to grow. In 1860, there were nine port cities that had populations over 100,000 (Boston, New York, Brooklyn, Philadelphia, Baltimore, New Orleans, Chicago, Cincinnati, and St. Louis)—by 1910, there were fifty. Just as significantly, the country's largest cities were no longer confined to the eastern seaboard or to the Midwest. There were several large cities in the plains states, and half the population of the far west was living not in its fertile valleys or at the feet of its glorious mountains, but in its cities: Los Angeles, Denver, San Francisco, Portland, and Seattle. By 1920, for the first time in the nation's history, just slightly over half of all Americans lived in communities that had more than 10,000 residents.

Money was flowing in the same direction that people were; by 1900, the nation's wealth was located in its cities, not in its countryside. The nation's largest businesses and its wealthiest individuals were in its cities. J. P. Morgan and Cornelius Vanderbilt controlled their railroad empires from New York; Leland Stanford and Charles Crocker ran theirs from San Francisco; John D. Rockefeller operated from Cleveland and New York; Andrew Carnegie, at least initially, from Pittsburgh. Probably by 1880, and certainly by 1890, stock exchanges and investment bankers had become more important to the nation's economic health than cotton wharves and landed gentry.

This transition to an urban society had political consequences because political power tends to follow the trail marked out by wealth (and, in a democracy, to some extent by population). In the early years of the nineteenth century, when the independent political character of the nation was being formed, most Americans still lived on farms and American politics was largely controlled by people who earned their living directly from the land. After the Civil War, city residents (being both more numerous and more wealthy) began to flex their political muscles and to express their political interests more successfully. The first twelve presidents of the United States had all been born into farming communities, but from 1865 until 1912, the Republican party, then the party that most clearly represented the interests of big business and cities, controlled the White House for all but eight years, and those eight years were the two terms served by Grover Cleveland, who before becoming president had been the mayor of Buffalo, New York.

The transition to an urban society also had economic and technological consequences. In a kind of historical feedback loop, industrialization caused cities to grow and the growth of cities stimulated more industrialization. Nineteenth-century cities were, to use the term favored by urban historians, walking cities. Since most residents could not afford either the cost or the space required to keep a horse and carriage, they had to be able to walk to work or to work in their own homes. Since businesses also had to be within walking distance of each other, this meant that as cities grew they became congested; more and more people had both to live and to work within the same relatively limited space. With congestion came disease; all nineteenth-century American cities were periodically struck by devastating epidemics: cholera, dysentery, typhoid fever.

Even before they understood the causes of these epidemics, city governments became convinced that they had to do something both to relieve the congestion and to control the diseases. Streets had to be paved, running water provided, sewers constructed, new housing encouraged. This meant that reservoirs had to be built, aqueducts and pumping stations constructed, trenches dug, pipes purchased, brickwork laid, new construction techniques explored. All of this municipal activity not only stimulated American industry but also served as a spur to the growth of civil engineering.

In addition, in the years between 1870 and 1920, many American cities actively stimulated industrialization by seeking out manufacturing interests and offering operating incentives to them. Many of the nation's older cities found themselves in economic trouble as railroad depots became more important than ports as nodes in the

country's transportation system. In their distress, these cities decided that their futures lay not in commerce but in manufacturing, and they began to seek out manufacturing entrepreneurs to encourage industrial growth. By that time, the steam engine having been perfected and its manufacture made relatively inexpensive, manufacturers had ceased to depend on waterwheels as a power source, which meant that they could easily (and profitably) establish their enterprises in cities rather than in the countryside; the development of the electric motor only served to increase this potential.

Minneapolis became a center of flour milling, Kansas City of meatpacking, Memphis of cotton seed oil production, Rochester of shoe manufacturing, Schenectady of electric equipment, New York of ready-made clothing, Pittsburgh of steel and glass manufacture. Local banks helped manufacturers start up in business and local politicians helped recruit a docile labor force, all in the interests of stabilizing or augmenting a city's economy. Nationwide, the net result was a positive impetus to the growth of industry; the processes of industrialization and urbanization are mutually reinforcing.

If American cities grew prodigiously during the second half of the nineteenth century, so, too, did the American population as a whole: between 1860 and 1920, the population of the United States more than tripled (from 31 million to 106 million). Some of the increase was the result of a high natural birthrate; in general, American families were larger than what is needed to keep a population at a stable size from one generation to the next. In addition, as the result of improvements in public health and improvements in the food supply, the death rate was declining and life expectancy was rising. People were living longer and that meant that in any given year a declining proportion of the total population was dying. On top of this, immigrants were arriving in record numbers. The figures are astounding; the total, between the end of the Civil War and the passage of the Immigration Restriction Acts (1924), came to over 30 million people. Like their native-born contemporaries, immigrants had a high birth rate and a declining death rate and more of their children lived past infancy and then enjoyed a longer life expectancy, all of which further contributed to the mushrooming size of the American population. This startling population increase—almost 20 percent—reflects another crucial difference between societies that have become industrialized and those that have not. In a preindustrialized society, the size of the population changes in a more or less cyclical fashion. If the weather cooperates and the crops are bounteous and peace prevails, people remain reasonably healthy and many children live past infancy; over the course of time the population will grow. But eventually the population will grow too large to be supported by the available land or the land itself will become infertile. Droughts may come or heavy rains; locusts may infest the fields or diseases may strike the cattle. Men will be drawn off to battle just when it is time to plow the fields or soldiers engaged in battles will trample the wheat and burn the barns. Then starvation will ensue. People will succumb to disease; fewer children will be born, and more of them will die in infancy. The population will shrink.

Under preindustrial conditions, such population cycles have been possible. Sometimes the cycle will take two generations to recur, sometimes two centuries, but it has recurred as long as there have been agricultural peoples who have been keeping records of themselves. Industrialization breaks this cyclical population pattern. Once a country has industrialized, natural disasters and wars do not seem to have a long-term effect on the size of its population; the rate of increase may slow for a few years or so, but there is still an increase. And the standard of living keeps rising as well. People stay relatively healthy; they live longer lives. Generally speaking, they can have as many (or as few) children as they want, knowing that, also generally speaking, most of their children will live past infancy. This is the salient characteristic that makes underdeveloped countries long for development: industrialized countries seem able to support extraordinarily large populations without any

long-term collapse either in the size of the population or in the standard of living.

Industrialized countries can do this because agriculture industrializes at the same time that manufacturing does. In the transition to industrialization, what is happening on the farm is just as important as what is happening in the factories since, to put it bluntly, people cannot work if they cannot eat. These social processes—sustained growth of the population and the industrialization of agriculture—are interlocked. Both were proceeding rapidly in the United States between the years 1870 and 1920 as American farmers simultaneously pushed west and industrialized, settling new territory and developing more productive farming techniques. As the frontier moved westward, roughly 400 million new acres were put under cultivation: virgin prairies became farms, fertile mountain valleys were planted in orchards, grassy hills became grazing land for sheep and cattle. The total quantity of improved acreage (meaning land that had been cleared or fenced or otherwise made suitable for agricultural use) in the United States multiplied two and a half times between 1860 and 1900.

This alone would have considerably expanded the nation's agricultural output, but newly introduced agricultural implements profoundly altered the work process of farming (particularly grain growing) and increased its productivity. The first of these was the reaper (patented by Cyrus McCormick in 1834 and in limited use even before the Civil War). The reaper, which was pulled by horses, replaced hand labor. Once a reaper had been purchased, a farm owner could quadruple the amount of acreage cut in one day or fire three-day laborers who had previously been employed for the harvest or greatly increase the acreage put to plow (since the number of acres planted had always been limited by what could be reaped in the two prime weeks of harvest).

The reaper was followed by the harvester (which made binding the grain easier), followed by the self-binder (which automatically bound the grain into shocks), and—in the far west—followed by the combine, a steam-driven tractor threshed and bagged the grain automatically, sometimes at the rate of three 150-pound bags a minute). In those same years, haymaking was altered by the introduction of automatic cutting and baling machinery, and plowing was made considerably easier by the invention of the steel plow (John Deere, 1837) and the chilled-iron plow (James Oliver, 1868), both of which had the advantage of being nonstick surfaces for the heavy, wet soils of the prairies.

The net result, by 1900, was that American farmers were vastly more productive than they had been in 1860. Productivity has two facets: it is a measure both of the commodities being produced and of the labor being used to produce them. Statistics on wheat production indicate how radically American agriculture was changing in the second half of the nineteenth century. In 1866, there were roughly 15.5 million acres devoted to wheat production in the United States; farmers achieved average yields of 9.9 bushels per acre, resulting in a total national production of about 152 million bushels. By 1898, acreage had roughly trebled (to 44 million), yields had almost doubled (to 15.3 bushels per acre), and the total production was 675 million bushels.

All this was accomplished with a marked saving of labor. By the hand method, 400 people and 200 oxen had to work ten hours a day to produce 20,000 bushels of wheat; by the machine method, only 6 people (and 36 horses) were required. Farms were getting larger, ownership was being restricted to a smaller and smaller number of people and more machinery was required for profitable farming (between 1860 and 1900, the annual value of farm implements manufactured in the United States went from $21 million to $101 million)—at the same time, the farms were becoming more productive.

What this means, put another way, was that a smaller proportion of the nation's people were needed to produce the food required by its ever larger population. Some people left their farms because they hated the farming life, some because they could not afford to buy land as prices began to rise, some because they

profitability of small farms. The farming population (this includes both owners and laborers) began to shrink in relation to the rest of the population.

New transportation facilities and new food-based industries made it easier and cheaper for the residents of cities and towns to eat a more varied diet. The fledgling canning industry was spurred by the need to supply food for troops during the Civil War. After the war, the canners turned to the civilian market, and by the 1880s, urban Americans had become accustomed to eating canned meat, condensed milk (invented by Gail Borden in 1856), canned peas, and canned corn. The Heinz company was already supplying bottled ketchup and factory pickles to a vast population, and the Campbell's company was just about to start marketing soups. By 1900, cheese and butter making had become largely a factory operation, made easier and cheaper by the invention of the centrifugal cream separator in 1879.

After the Civil War, the railroads replaced steamboats and canal barges as the principal carriers of farm products (from wheat to hogs, from apples to tobacco), thus both shortening the time required to bring goods to market and sharply lowering the cost of transportation. After the 1880s, when refrigerated transport of various kinds was introduced, this trend accelerated: even more products could be brought to market (butchered meat, for example, or fresh fish) in an even shorter time. New refrigeration techniques transformed beer making from a home to a factory operation; by 1873, there were some 4,000 breweries in the United States with an output of 10 million barrels a year. Commercial baking had also expanded and Americans were becoming fond of factory-made crackers and cookies. In the end, then, another historical feedback loop had been established, a loop connecting industrialization with agricultural change. Industrialization made farming more productive, which made it possible for the population to increase, which created a larger market for manufactured goods, which increased the rate of industrialization.

CONCLUSION: INDUSTRIALIZATION AND TECHNOLOGICAL SYSTEMS

By 1920, a majority of Americans had crossed the great divide between preindustrial and industrial societies. The foods they ate, the conditions under which they worked, the places in which they lived—all had been transformed. The majority of Americans were no longer living on farms. They were eating food that had been carried to them by one technological system (the railroad) after having been processed by machines that were powered by a second (electricity) and lubricated by a third (petroleum). If they wanted to light their domiciles at night or heat their dwelling places during cold weather, they could not avoid interacting with one or another technological system for distributing energy—unless they were willing to manufacture their own candles (even then, they might have ended up buying paraffin from Standard Oil). The social ties that bound individuals and communities together—someone has been elected, someone else has died, young men are about to be drafted, a young woman has given birth—were being carried over, communicated through, and to some extent controlled by technological networks that were owned by large, monopolistically inclined corporations. More people were living longer lives; fewer babies were dying in infancy; the standard of living for many Americans (albeit not for all) was rising. And at the very same time, because of the very same processes, people were becoming more dependent on each other.

Early in the nineteenth century the process of industrialization had appeared (to those who were paying attention) as a rather discrete undertaking: a spinning factory in a neighboring town, a merchant miller up the river, a railroad station a few miles distant. By the end of the century, virtually all Americans must have been aware that it had become something vastly different: a systematic undertaking that had created interlocking physical and social networks in which all Americans—rich or poor, young or old, urban or rural—were increasingly enmeshed.

1.1.3 CONTAINERIZATION AND GLOBALIZATION: AN INNOVATION AND ITS IMPACT

ARTHUR DONOVAN

The need to move goods from one location to another dates far back into antiquity. Whether items were being carried in a basket, slung from the back of a camel, or piled on ox carts, it is safe to assume that there was always an interest in moving goods more efficiently, either to save steps or to increase profits from trade. Many improvements were made in the ways in which goods were transported, but none had the impact of the approach in the title of this selection—containerization.

Technological advances played a critical role at all stages of the evolution of what today is called "globalization," from the earliest stages to the sophisticated levels with which we are now familiar. With far-flung factories in less-developed countries making much of what is consumed in the developed world today, it is essential that the output of those operations be moved quickly, cheaply, and safely to distant markets.

Although some goods are moved by air, the far more cost-effective way to do this is by sea. The steam engine, initially developed to pump water from deep mines in England, was improved and adapted to powering factories, ships, and trains. Along with enhanced navigational technologies including chronometers and more accurate compasses, shipping took on greater importance in world commerce. Up until the period discussed in this selection, all shipped goods were loaded onto ships by dock workers, stacked in the holds of those ships, and, when the ship arrived at its destination, unloaded in the same way but in reverse order.

Although transportation of people and goods by water was seen as more efficient than over-land travel, the manner of loading and unloading ships increased labor costs and made pilferage relatively easy. If world trade was going to grow, a better system had to be developed.

As with many innovations, it is likely someone else would have eventually put together the various pieces into a practical system had Malcolm McLean not done so in 1956. All the necessary pieces were available to many and needed only to be brought together in a unique and insightful way. McLean did that and the rest is, as they say, history.

Certainly, not everyone saw this change as beneficial. Workers whose jobs and livelihoods were threatened by containerization were less than enthralled with those metal boxes. Although some new jobs were created in manufacturing both by the containers and other special equipment required to handle them, many jobs formerly done by well-paid, unionized stevedores were eliminated and would never come back.

Source: From *Science, Technology, and the Humanities: A New Synthesis* pp. 67–74, Copyright © 2011 by Steven Institute of Technology. Reprinted by permission of Arthur Donovan.

FOCUS QUESTIONS

1. What were the key technological innovations that led to the widespread use of containers in global trade? How might the myriad of products available to shoppers be different today if containers were not used to the extent they are?
2. How do you think you would have reacted to the increased utilization of containerized shipping if you were a stevedore? What if you were involved in some other aspect of the business world?
3. What other factors might have contributed to the appeal containerized shipping and helped it become dominant in a relatively short time? What might have happened to McLean's idea if those conditions had not existed?
4. What role has containerization played in the broad area of globalization, and in what ways has this phenomena both negatively and positively impacted the lives of many Americans?

KEYWORDS

break-bulk, burgeoning bourgeoisie, Commerce Commission, container (ization), core technology, economic globalization, globalization, intermodalism, Interstate integrated freight system, long-haul freight, standard-sized box

On April 26, 1956, an important event in the history of freight transportation took place in Newark, New Jersey. The central figure in this event was Malcom McLean, a North Carolina businessman who during the previous two decades had expanded his trucking company from a one-man start-up into one of America's largest and most successful interstate haulers. By 1953 McLean Trucking was running 1,425 trucks and trailers up and down the east coast, had 37 terminals from Maine to Florida, and was annually grossing over nineteen million dollars. But McLean was not yet satisfied, for he knew all-too-well that the highways his trucks depended on were heavily congested and were being pounded into powder by increasing traffic. Pressed to find a better way to move long-haul freight, McLean in 1954 asked the Interstate Commerce Commission for permission to begin moving his truck-trailers by sea. This initiative soon resulted in the development of a technological innovation of historic importance, a new freight transportation system that was given two different names. The standard-sized box McLean developed to move general cargo came to be called a container, and his new system for moving these boxes is now known worldwide as "containerization." The other name given to this new system, "intermodalism," refers to the fact that the use of containers makes possible the mechanization and efficient integration of the three major modes of surface freight transportation: road, rail and water. But whichever name is used, it is now universally acknowledged that in 1956 McLean's introduction of containers launched a worldwide revolution in the way general freight is packed and moved over land and sea.[1]

McLean was an ambitious entrepreneur who was always looking for new ways to expand his trucking company and capture more business. When he was ready to show his new container system to potential customers, he invited representatives from the shipping industry to a dockside reception at the headquarters of the Pan-Atlantic Steamship Company, a coastwise shipping group he had recently acquired. The guests who gathered on the waterfront on that blustery April day included reporters from the shipping press, officials from the Port Authority of New York and New Jersey, and shipping agents from several companies that moved goods in and out of the New York region. As his guests were enjoying their sandwiches and salads, McLean directed their attention to a nearby crane that was hoisting truck-trailer-sized cargo boxes onto a special deck that had recently been installed on a World War II tanker that was lying alongside a nearby pier. When the new deck had been fully loaded with fifty-eight cargo boxes, the ship, given the unusual name *IdealX*, departed Port Newark, bound for Houston, Texas. Ever since that day, the departure of the *IdealX* has been remembered as the first day in the container

revolution, and Malcom McLean has been deservedly acknowledged as the entrepreneur who launched modern container shipping.

The new kind of commercial freight service that McLean showed his guests in 1956 soon attracted considerable attention, and despite the traditional shipping industry's deeply embedded resistance to change, in just a single generation containerization transformed the way general freight is moved over land and sea throughout the world. And as we look back at this turning point in the history of transportation, it is clear that McLean's willingness to walk away from a successful trucking company to start this new venture says a lot about him. He was a bold entrepreneur who was convinced that the cost of moving general freight could be dramatically reduced by radically reorganizing how it is packed and handled. Instead of moving each bag, barrel, and box of cargo piece-by-piece when shifting freight from one mode of transportation to another, McLean packed cargo into sturdy truck-trailer-sized boxes that were sealed shut before being loaded onto ships specially fitted out to carry them to distant ports. These boxes, which soon became universally known as containers, could also be moved by rail by putting them on railroad flatcars, or over highways by attaching them to trailer chassis. For McLean, containerization was simply an extension of trucking by other means, and as everyone in the freight business knew, cargo-handling costs aside, it was far less expensive per mile to move cargo by ship or by train than by truck. If McLean's experiment was successful, the reduced costs made possible by eliminating the repeated handling of individual parcels of freight would enable him to capture far more long-haul freight business than he could possibly have captured by using trucks alone.

McLean did not invite his guests to lunch that day because he wanted them to invest in his new company. He wanted their freight, not their capital. He wanted them to abandon the traditional break-bulk system of cargo handling and pack their goods in his intermodal containers instead. He was convinced that the cost advantages provided by Sea-Land Service, the name he gave his new company, would enable him to capture much of the freight that had previously moved up and down the Atlantic coast by rail, truck, offering a new kind of integrated freight service, a service that soon came to be called intermodalism, and he urged his guests to watch carefully as his first cargo boxes, each of which, he assured them, was loaded with paying freight, were hoisted aboard *IdealX*. The luncheon he hosted that day did what he hoped it would do, and Sea-Land Service soon had all the business it could handle.

April 26, 1956, was Day One of the container revolution, and Malcom McLean is still rightly remembered as the man who put boxes on ships.[2] Before containerization, as old-timers still recall, break-bulk cargo-liners would stay in port alongside a dock for up to two weeks while longshoremen off-loaded incoming cargo and then loaded and stowed out-going cargo. Today, however, the break-bulk cargo-liners of old and the piers they docked at during their extended stays are long gone. Gone too is the infamous pre-container waterfront world of heavy lifting, theft and blue-collar crime, a world explored dramatically in the classic film On the Waterfront. Today's containerports are huge, heavily-mechanized outdoor factories, into which enormous containerships loaded with thousands of boxes glide with elephantine grace. Once firmly secured alongside their designated dock, the ship's hatches are opened and looming container cranes begin hoisting boxes out of the holds. In a matter of hours the containers destined for that port have been unloaded, outbound containers have been loaded aboard, and the ship is ready to depart. The old notion that sailors go to sea to see the world no longer applies, for the pace of transportation in the container age is relentless.

This new way of packing and moving general cargo expanded rapidly during the more than fifty years that have passed since McLean planted the seed of containerization in Newark, New Jersey. During this half-century the adoption of containerization around the world has transformed both seafaring and the port cities that once drew their wealth and vitality from their working waterfronts. Containerization is, therefore, much more than just one industry's new technology. It launched and sustains a revolution that has affected the lives of all those who work in industries that are in any way linked to international trade in both the highly developed and the developing worlds. Of course McLean

innovation would have; he was just looking for a better way to move freight. But within his lifetime the new system he developed for packing and moving general cargo did indeed change the world.

The story of containerization provides, I believe, an especially revealing example of how a specific innovation can, given the right circumstances, launch a global revolution, the innovation in this case being one that transformed the way general freight is moved over land and sea. But what have the consequences of this revolution been? Perhaps the best way to engage this question is to ask what we now mean when we speak of "globalization" and "the global economy." It is no accident that contemporary economic globalization emerged during the first generation of containerization. But what are the linkages between this innovation in freight transportation and the more comprehensive economic revolution called globalization, the new level of economic integration in which containerization now plays such a prominent role? The answer is not a mystery, but neither is it immediately obvious.

When discussing this question with students, I like to begin with a simple and quite concrete question. I ask them what kinds of choices are available to them when they are shopping for beverages at their local Beer Barn. For me this is a Socratic moment, an opportunity to pose a simple question that reveals more than is expected about the world at large. So I ask my students if their local Beer Barn sells beer brewed and bottled outside the United States. If it does, and they all now do, I then ask them to explain how beer that was made in Bangkok or Amsterdam or Australia can compete in price with brands that are produced only a few miles away from the store in which it is sold? What about transportation costs, the cost of distance? And once you start asking such questions, it becomes obvious how astonishingly novel today's global economy really is. Quite a large lesson can then be teased out of this utterly mundane point-of-entry.

To see more clearly the linkages between containerization and globalization, one can also begin with Wal-Mart. Today it is widely known that Wal-Mart is the world's largest store, but you may not know that Wal-Mart, along with Home Depot, Target, and a long list of other Big-Box stores, is overseas, goods that are packed in containers and sent to stores all across the United States and other countries around the world as well. The connections between containerization and globalization are important, for container shipping has transformed how international commerce is conducted today. The wide spread use of containers and the recent expansion of the global economy have proceeded hand-in-hand. Of course other technical advances have contributed to these developments as well, for containerization alone would not have been enough. And we must also keep in mind that global commerce today is vulnerable to sudden interruptions, as the 9/11 attacks in New York City so tragically demonstrated. The world of commerce has always been and still is subject to piracy and other forms of raiding, and trade must, as always, be protected. Containerization can also be used for other nefarious purposes as well, be it smuggling or sabotage. Yet the advantages containerization provides when used in the service of legitimate exchange are so great that occasional misuses have not yet seriously impeded the new technology's rapid growth. Simply shutting down global container shipping to avoid terrorist attacks is not currently considered an option, which is why the Department of Homeland Security is spending so many of your tax dollars on policing the container trades. But the basic problem remains, and only time will tell if in the long run the benefits of containerization will outweigh its associated vulnerabilities. I don't think we should be pessimistic about this. The use of containers, like driving on crowded highways at break-neck speeds and so much else in modern life, involves risks, but so far these risks have been deemed acceptable and have not slowed containerization's rapid growth.

Containerization is now a core technology in today's global economy, and today's global economy reaches further and involves a greater portion of the entire world than any of the imperial or commercial empires that humankind managed to construct in earlier eras. And as economists from Adam Smith to the present have pointed out, trade and the specialization it encourages are great generators of wealth. Even a brief look at any of today's great containerports reveals that containers carry a vast volume of highly-valuable trade. Containerization

production and marketing of all sorts of manufactured goods. It has enormously reduced the costs and extended the range of global freight transportation, making it possible to treat nearly any location on the entire globe as both a potential center of production and a market for goods produced elsewhere.

But for those of us who study the origins and consequences of science and technology, simply knowing what containers are and how they are used is not enough. We also want to know how this new technology was created, how it grew into a global system, and how it has transformed world commerce. We also want to know how this new system operates and how its growth has changed the way people now live. Simply put, we need to understand both how containerization came to be and what, in human terms, it means. Containerization, a relatively new topic on the historical horizon, poses interesting challenges for historians of technology, and especially those who study transportation and globalization.

But, you may say, haven't we been through all this before? Didn't Marshall McLuhan point out that we all now live in a global village, a universe of shared information made possible by modern communications? Yes, that's true, and McLuhan was indeed an original thinker who focused on a particular set of modern technologies, "the media," and explored the consequences of their recent extension to global proportions. And it is also true that without the worldwide communication systems made possible by computerization and the internet, today's container industry would not have the information-rich "nervous system" it needs to maintain "command and control" over the movement of containers around the world. Perhaps then the development of modern media and McLuhan's Global Village should be thought of as steps toward the development of containerization. But we also need to remember that Marshall McLuhan and Malcom McLean were addressing very different subjects. McLuhan, a lively-minded professor of cultural studies, and McLean, a plain-spoken trucker whose formal education ended with high-school, saw the world through very different lenses. What they had in common was a lively-minded interest in how recently-developed technologies could shape the future.

For McLean, the main challenge was to minimize the cost of freight transportation, and he did this by using standard-sized cargo boxes that can be moved over roads, rails, and water without requiring that the goods packed in them be repeatedly unloaded and reloaded when shifting from one mode of transportation to another. Of course McLean could not foresee all the consequences of his innovation, but he nevertheless seized upon the idea of containerization and used it in ways that produced a new kind of global economy. And make no mistake about it: the global economy in the age of containerization is astonishingly new. Consider, for instance, that today it costs only 83 cents per item to load 4,000 video-cameras, all packaged and ready to be put on store shelves, into a single 40-foot container and transport them from Asia to Europe. Similarly, whiskey distilled and bottled in Scotland is shipped in containers to Japan at a transportation cost of only 4.7 cents per bottle.[3] These figures, astonishing as they are, are in no way exceptional, for today the cost of moving goods from where they are manufactured to the markets in which they are sold has become almost negligible. In this sense, the lesson learned in the Beer Barn is enlightening, for we now really do live in a different world, one in which places of production and market locations can be both highly dispersed and globally integrated, a world in which the cost of distance has become almost negligible.

Consider, for instance, the story of Zhang Yin, the "Queen of Trash," now the wealthiest woman in China and one of the richest self-made women in the world.[4] Her father, a military officer, was imprisoned during the Cultural Revolution of the 1960s and was later rehabilitated. In the 1980s, when China began experimenting with capitalism, she moved to Shenzhen, a city on China's southern coast, and began working for a joint-venture trading company that supplied paper to factories in Hong Kong. She soon found that market too limited and moved to Los Angeles, where she and her husband formed a company that bought scrap metal and paper in America and sold them to businesses in China. These materials were carried to China at very low rates in containers that were being returned after delivering manufactured goods of much higher value to the United States. If these boxes had not been carrying recycled paper and scrap metal to China, they would have been returned empty.

It was an ideal arrangement: low cost materials gathered in resource-rich America and low cost shipping to China, where the demand and prices for these recycled materials were much higher than their cost. Zhang's Nine Dragon Paper Company is now the biggest papermaker in China, and *Forbes Magazine* recently listed her as China's fifth richest person.

Zhang's story is dramatic, but not unique, for containerization has created many successful new capitalists by linking, at very low cost, distant suppliers and new groups of consumers. Recycled paper is not the only product that is sent to Asia in containers that are being recycled after bringing to America high-value consumer goods manufactured in Asia. Hay used to feed cattle is another bulk product exported in this way. The hay is grown and harvested on the high plains of the American West and is then compressed and packed in containers that are being recycled to Japan and other Asian nations. The hay is then fed to beef cattle that in turn supply the growing East Asian market for meat. Of course it comes as a bit of a surprise to hear that waste paper and compressed hay are two of America's leading exports to Asia, but in reality this exchange is nothing more than straightforward market exchange at work. But it is international commerce as it now operates in the age of containerization, and that makes all the difference.

Where is all this leading? Well, the weekly newsmagazine *The Economist* recently argued that despite all the problems plaguing the global economy, there are grounds for optimism. In a special section on "the new middle classes in emerging markets,"[5] the magazine focused on the economic gains of leading emerging nations that are benefiting from globalization and the growth of trade. The groups that are pressing forward, *The Economist* argues, are a "burgeoning bourgeoisie," a new "middle-class." Four key factors are cited to support this view:

1. "The middle class matters because it does a lot of consuming." (p. 14)
2. "Education is high on any list of middle-class values." (p. 16)
3. "The middle class's third distinctive contribution to growth is its gift for entrepreneurship." (p. 17)
4. "The middle class's [fourth] distinctive contribution to growth [is] its support for the right economic policies and its backing for democracy." (p. 17)

Well, that's quite a bold set of assertions and conclusions. Yet it does seem reasonable to suppose that contemporary globalization, which to a large extent is made possible by containerization, is creating new and realizable opportunities for those who were kept on the margins of economic development in earlier regimes and during the Cold War. And these opportunities well may enable peoples who have long been kept at a distance to now find a place at the global table. Of course not all nations are now on the road to greater trade and increasing wealth, but when they are prepared to follow it, the path they need to follow has been marked out for them. And perhaps at this moment in history container shipping will assist in the creation of a more equitable and truly global economy. And perhaps then many more people will be able to enjoy beer and wine and other consumer goods from around the world as well, goods purchased at low prices that those who live in developed countries now simply take for granted.

NOTES

1. Arthur Donovan and Joseph Bonney, *The Box That Changed the World: An Illustrated History* (Commonwealth Business Media, East Windsor, New Jersey, 2006), p. 35. The American container industry commissioned this history as part it its celebration of the first fifty years of containerization. The book is not available from bookstores, but can be purchased from the Containerization & Intermodal Institute (containerization.org), or the publisher, (cbiznedia.com). Current plans are to include this book, hereafter cited as TBTCTW, in the forthcoming on-line library being developed by Google.
2. The history of containerization, like the industry itself, has expanded rapidly. The industry first and most profoundly transformed by the new cargo-handling technology was liner shipping, and one of the first books, and still one of the best, on this subject is by Frank Broeze: *The Globalization of the Oceans: Containerization from the 1950s to the Present*, (International Maritime Economic History Association, Research in Maritime History No.23, St. John's, Newfoundland, 2002). A more recent

treatment that integrates containerization into the long history of ocean shipping is Martin Stopford's masterful *Maritime Shipping*, 3rd ed. (Routledge, 2009). The best-known popular account of the container revolution is by an economist, Marc Levinson, *The Box* (Princeton University Press, 2006).

3. Martine Stopford, *Maritime Economics*, 3rd ed., (Routledge, 2009), p. 512.
4. *International Herald Tribune*, 15 Jan 2007; see also Evan Osnos, "Wastepaper Queen," *The New Yorker*, 30, March 2009, pp. 46ff.
5. Issue dated February, 2009, pp. 14–20.

1.1.4 HISTORY AS TECHNOLOGICAL CHANGE

ROSALIND WILLIAMS

In this selection from Williams's book *Retooling: A Historian Confronts Technological Change*, MIT's dean of undergraduate education and student affairs reflects on how the idea of historical change has become nearly synonymous with technological change and how the term *technology* has been reduced to meaning only information technology. While acknowledging the important role that technological innovation has played in human history, Williams wants to resist the idea that all historical change can be reduced to or explained by technological change. But she also observes that we are now on the verge of another great transformation of the human habitat, comparable in significance to those brought about by the agricultural revolution and the Industrial Revolution. The question is, how will this transformation affect our sense of human identity? Where will we be able to find meaning and purpose in a world that is our own creation?

FOCUS QUESTIONS

1. What does Williams mean when she says that "technological innovation resembles but hollows out the idea of progress"?
2. What are some examples of previous technological revolutions, and how have they transformed the human habitat? According to Williams, what kind of transformation of the human habitat appears to be taking place because of the information revolution?
3. What does she mean when she describes the world of the information age as a reflexive, or "boomerang," world?
4. Discuss and evaluate her claim that "history and nature are weakening as well-defined external frameworks that give meaning to human life."

KEYWORDS

artificial habitats, information technology, innovation, progress, reflexivity, technological revolutions

One dominant assumption about the information age is that it represents a "new world" of technological change that relentlessly pushes into retreat an "old world" of culturally driven resistance, much as the soldiers of the Roman Empire pushed the Celts back into the western hills of Britain. Today's conflict, we are told, is between the new economy and the old one, between digital and analog, between

Source: Rosalind Williams, "History as *Technological Change*" from *Retooling: A Historian Confronts Technological Change*, pp. 14–26, © 2002 Massachusetts Institute of Technology, by permission of The MIT Press.

systems and life-world, between globalization and identity. The human race is split between dynamic "change agents" and the stagnant nonretooled, between neo-nomadic "symbolic analysts" who work with bits and a "disposable labor force that can be automated and/or hired/fired/offshored...."[1]

Such binary thinking is extended to history itself. The information revolution is leading us into a new phase of history, in which technological change will be the major and constant determinant of human life. History becomes the record of a grand struggle between the irresistible force of technology-driven "change" or "innovation" on the one hand and misguided if understandable culturally driven resistance to change on the other, with "change" and "innovation" inevitably winning.

This is supposedly an age without a sense of history, but the view of history as a struggle between technological change and technological resistance is powerful and pervasive today. Implicit rather than explicit, it is less a theory (which would imply some conscious awareness of the issues and some effort at logical consistency) than an ideology (ideologies are not especially logical or consistent).[2] It is an ideology of history in which the very word "history" has been displaced by "technology." Instead of being a figure in the ground of history, technology has become the ground—not an element of historical change, but the thing itself. We have come to assume that where technology is going is where history is going, as if they are now one and the same.

The ideology of technological change has become the dominant, if implicit, theory of history in our time, even for those who never pick up a history book. In Karl Mannheim's famous distinction, ideology is a set of "wish-images" intended to evolve the existing order into a better one, while Utopia represents a wish to break the bonds of the existing order, to overturn it for a new reality. Mannheim adds: "The Utopia of the ascendant bourgeoisie was the idea of 'freedom.'"[3] Innovation has come to define a Utopia of freedom, an ideal place running by market and technological laws, where creative minds and free-flowing capital unite to make a new world abundant in possibilities and energy. This is a much more appealing world than the one we actually live in, where innovations become hardened into bureaucracies, where actions have consequences, and where waves of change wash back on shore. The Utopian "no place" of market-driven innovation presents an agreeable escape from the complexities and consequences of a crowded, risky world. Mannheim emphasizes that ideology and Utopia are ideal types; in practice, it is often difficult to distinguish them within the social order or even within an individual's mind. In the view of history as technological innovation, however, we can often discern a set of "wish-images" that is Utopian without being especially progressive.

As an MIT administrator, I rarely heard anyone talk explicitly about history, but I often heard people express the implicit view of history as technological change. They did this through repeated use of a cluster of deceptively simple words: "technology," "change," "culture," "space," "time," "community." These words are highly interdependent, deriving their meaning from the others and forming a closed world of cross-references. They are also highly reflexive, acquiring meaning from the very changes they are supposed to analyze.[4] They also drain meaning away from other, formerly more significant words, such as "engineering" and "progress."

The sun of this linguistic system is "technology." In 1861, when MIT was given its name, the word was novel, even daring, and grandly comprehensive. Proclaimed in chiseled stone below the Great Dome, "technology" still defines, in a lofty way, the Institute's mission. But in everyday meetings below and around the dome "technology" has little grandeur. It has come to mean *information technology*—a recent, drastic, telling reduction. In committee meetings, when we sketched business processes or organizational models on a flip chart, someone would write "technology" with a circle around it. We would talk about designing new learning spaces to accommodate "the technology." We would write and review reports saying things like "The primary motivation to revise the policy [is] the recognition that technology has transformed the means by which we collect, manage, and provide access to student information."

At no time, in any of these contexts, did anyone suggest that "technology" might include

more than information technology, nor did anyone challenge its status as an independent entity. "Technology" once was a grand term, though also a dangerous one, because it was so often used as an abstract, independent historical agent.[5] It is still dangerous in this way, but it is no longer grand. It means information technology and it means change.

While "technology" expands its rhetorical reach, that of "engineering" shrinks. Never a glamorous term (though it was a solid one), it is now used more rarely, and mostly in a connection with a specific project or department. Even more conspicuous in its absence is the word "progress." Not long ago, discussions of technology were dominated by "progress talk."[6] At MIT today, people talk a lot about scientific progress, referring to the kind of discovery that would win a Nobel Prize, but little about technological progress, except with irony.

Instead, "technology" is linked with "change." Everyone talks about change as relentless and inevitable but also, almost always, as pointless, in the double sense of lacking a purpose and lacking an end. Progress has a story line; change does not. After one meeting about new software systems, a participant told me that a consultant had advised her to use the concept of a "change journey" to help people overcome their resistance to change. When I inquired what was the destination of such a journey, I was told that it might be thought of as arriving at a point of comfort with change. "Change," like "technology," has been sharply reduced in meaning. Of all the possible types of change—political, intellectual, cultural, social—technological change is all that really matters.[7] The change journey takes us through familiar landscapes of science, engineering, politics, the market, work, family. These too may change, everyone acknowledges, but the assumption is that they will do so only as a consequence of technological change.

Technological innovation resembles but hollows out the idea of progress.[8] Progress is measured in reference to human-defined goals: change and innovation are measured by market success. Market-driven innovation may tend to move in certain directions (Moore's Law, globalism, systems integration), but these directions are defined less by human purposes than by the rules of the market and the inner workings of technology itself. Innovation does not necessarily translate into human progress; any more than biological evolution means progress of the species (as Darwin well understood). Today's ideology of history as technological innovation turns the market, the motor of capitalist acquisition, into the motor of history itself, with technology as its agent.

It is easy to unmask the rhetoric of innovation and change, and to show that it is a cynical management tool justifying any sort of human disruption (moving, working harder or longer, being laid off) as the inevitable downside of the larger good. But then why is this view of history so widely, if semi-consciously, accepted?[9] Management babble is heavily discounted by most of us. In the midst of its ideological torrent, we try to keep living reasonably human lives, trying to figure out who we are, how to connect with others, how to accomplish something meaningful, how to enjoy the time we have. Still, the rhetoric of change management has some ring of truth, because it expresses, however inadequately, everyday experience in a world dominated by technologies we have created.

The rhetoric of change management has the effect, however, of trivializing the whole concept of change. Significant change is much more than a new software system or product. It is historical, in that it involves human relationships, expectations, and meanings. There is no single digital divide that we will step over collectively into the information age, leaving behind the creaky old world of everyday life for a sleek new world of frictionless systems. There are many fault lines in history today, and many of them run right through us.

The change agent and the change re-sister are often the same person.[10] We are living through a transformation of the human habitat that is much more comprehensive than what is usually thought of as "technological change."

THE NEW HUMAN HABITAT

Predictions about the information revolution have become the "airport reading" of the world—a substitute for beach reading in a harried age, just as predictable but a lot less engaging, celebrating not the steamy pleasures of physical reality but the

disembodied pleasures of virtual reality. Computing will be ubiquitous, our rooms will read our thoughts, the air will hum with messages, the Internet will become an omninet, we will be lifted out of the mud of localism to digital globalism—all on the other side of the great historical divide. Many historians of technology dismiss most of these predictions as unfounded digital evangelism. Historians are trained to look for continuity, to be wary of revolutionary claims, to spot inconsistencies. On the latter point, Alex Roland, a recent past president of the Society for the History of Technology (SHOT), notes that people who talk about technological revolution are often talking about different things: "Fifteen years ago we were confidently told that we were in a computer revolution. Now it is an information revolution. A more proper label might be a communications revolution. Or it may turn out in hindsight to have been a solid-state-physics revolution, or a digital revolution, or a micro-electronics revolution, or a control revolution. Some people think that the coming revolution in genetic engineering will outweigh all of these."[11]

So what kind of technological change is it that we are living through, and how does it compare with other major transformations in history? Is this another industrial revolution (the productive one of steam engines, cotton mills, and railroads), or is it more like the second industrial revolution (the consumer one of electricity, automobiles, cinema, and chemical industries)? Is the telegraph "the Victorian Internet,"[12] or is the Internet more like electricity? Or should we compare information to energy?

No analogy seems quite right. The industrial revolution of the late eighteenth and the early nineteenth century was defined by an unprecedented leap in human productivity brought about by the exploitation of the vast energy reserves stored in fossil fuels.[13] If you compare the efficiency of the steam-powered thresher in Laurel with the labor of hand threshing, and if you multiply that kind of efficiency millions of times in millions of settings, you begin to grasp the essence of the industrial revolution.

The second industrial revolution (which occurred in the late nineteenth century) also increased productivity, but its distinctive contributions were the distribution of cheap and flexible energy in the form of electricity, the construction of grid upon grid of communication and transportation systems (subways, automobiles, telephone, radio), and what contemporaries called "the democratization of luxury" through the proliferation of cheaper and more varied consumer goods. Collectively these technological innovations made everyday life far more comfortable and interesting than ever before. In the words of another past president of SHOT, Robert Post, "The late nineteenth century was the greatest period of technological change in terms of things that affected huge numbers of people's lives in basic ways—much more so than today."[14]

In all these discussions, historians are trying to develop a sort of historical Richter scale: What are the minor shocks, and what are the major ones? All of us are trying to get a sense of the scale of change in the midst of constant tremors. At one MIT meeting about educational technologies, a participant from the software industry commented: "Education is the building and technology is the earthquake." On the historical Richter scale, however, many historians of technology would say that the so-called information revolution is not so big. They disagree with the hypothesis of fundamental change preached by digital evangelists, in part because historians are used to measuring change by productivity rates and living standards. These measures miss some distinctive features of information technology. At the heart of information technology is the manipulation of symbols, not of matter. It does not simply "impact" culture and society; the technology itself is inherently cultural and social from the start. Directly as well as indirectly, it affects human experiences of space, time, communication, and consciousness.

From this perspective, the most appropriate historical analogy is another communications revolution: the invention of the mass media, or the printing press, or perspective, or numerals and the alphabet, or, behind them all, the invention of language—by far the most complex and sophisticated human invention, and a prerequisite for storing and passing on all other forms of technological knowledge and practice.[15] But this analogy

too is inadequate, insofar as information technology has transformed production and consumption as well as consciousness and communications.

If comparisons are hard to find, it is not because technological change is less significant today than in the past but because we are breaking it up into episodes instead of seeing it whole. The digital divide is a major episode, not the whole story. The technological revolution we are living through includes everything Roland names: computers, information, communications, solid-state physics, digital electronics, control, genetics. It also includes the two industrial revolutions, which now may be seen as partial and preparatory. The essential feature of the larger technological revolution is the creation of a new habitat for human existence.[16] On the historical Richter scale, it is The Big One. Elting Morison, a historian who joined the MIT faculty in 1946, described it as follows:

> We are well on the way, in our timeless effort to bring the natural environment under control, to replacing it by an artificial environment of our own contriving. This special environment has a structure, a set of tempos, and a series of dynamic reactions that are not always nicely scaled to human responses. The interesting question seems to be whether man, having succeeded after all these years in bringing so much of the natural environment under his control, can now manage the imposing system he has created for the specific purpose of enabling him to manage his natural environment.[17]

Morison's formulation shows how difficult it is to find adequate and precise (not to mention gender-neutral) language to describe the new human habitat. Calling our environment "artificial" or "special," as Morison does, suggests something unnatural or even abnormal about it. There is nothing more natural, however, than the human creativity and ingenuity that have shaped this new environment. For human beings, the production of technology is as natural as is the production of enzymes for certain bacteria.[18] Nor is an "artificial" environment "replacing" the natural one. Nature has ended not in the sense of going away but only in the sense of being so mixed and mingled with human processes that it can no longer be identified as a separate entity.[19] In the words of the social theorist Ulrich Beck: "Not a hair or a crumb of it is still 'natural,' if 'natural' means nature being left to itself."[20]

Human beings have always tried to control nature so as to make life safer, more predictable, more abundant, and more fulfilling. But since the beginning of recorded history, and back into unrecorded history, non-human nature has been the ground of human life. This relationship between technology and nature, between figure and ground, which had been reversing slowly over centuries, reversed decisively in the past century. The built world has become the ground of human existence, now framing and embedding non-human nature. We have gone from using technology to control and exploit our habitat to using it to detach ourselves from our habitat.[21]

We continue to struggle for words with which to express this new phase of history. As early as the first industrial revolution, Georg W. F. Hegel contrasted "first" or given nature with "second nature" created by humankind. In similar language, the historian of technology Thomas Parke Hughes has referred to the technological world as the "second creation." In the 1960s, Lewis Mumford contrasted "the organic habitat" with "this new megatechnics ... a uniform, all-enveloping structure, designed for automatic operation."[22] I find myself using the verbal shorthand of referring to a technological world, or, more accurately, a *hybrid* world—a world in which technology and nature are inextricably mixed.

There is a convincing historical precedent for the replacement of one human habitat by another: the Neolithic revolution, usually dated around 10,000 BCE and usually defined by the integrated invention of agriculture, of cities, and indeed of history itself, since the concomitant invention of writing permitted, for the first time, a durable collective memory. Left behind, except in specific settings and in deepest memory, was the pastoral life of hunting and gathering, which had been the human, or humanoid, way of life for hundreds of thousands of years.[23] All the transformative events of the Neolithic revolution—writing, agriculture, settlements—were based in material change, but they also redefined what it meant to be human.[24] As the historian of religion Mircea

Eliade notes, when most of humankind started raising crops "an ancient world, the world of nomadic hunters, with its religions, its myths, its moral conceptions, was ebbing away." Eliade continues, "Thousands and thousands of years were to elapse before the final lamentations of the old world died away, forever doomed by the advent of agriculture. One must also suppose that the profound spiritual crisis aroused by man's decision to call a halt and bind himself to the soil must have taken many hundreds of years to become completely integrated."[25]

The technological revolution of our times is the "decision"—collective, unconscious, incremental—to unbind ourselves from the soil. Eliade reminds us how far-reaching are the implications of binding ourselves instead to a self-constructed world. For millennia, the fact of settlement—humans living with other humans in a place over time—has shaped our ideas and practices of work, family, time and space, and society. This fact has also shaped our souls: solace, redemption, transcendence, and meaning have been defined by the twin powers of human community and nonhuman nature. It may be centuries before the lamentations for a lost world fade away and the human imagination can come to terms with its new habitat.

A technological world is partly natural and intensely human. The new fact of history on a material level is that nature is now so thoroughly intermixed with technology. The new fact of history on a social level is that we keep running into ourselves, as it were, as we build our values and our social order into the world. Sociologists describe this as the transition from modern society to something beyond: post-modernism, or (as Anthony Giddens prefers) post-modernity or radical modernity, or (as Ulrich Beck prefers) a risk society. Such thinkers stress the "reflexivity" of our society, meaning that we live in a world of echoes, a "boomerang" world (Beck's term) where everything that goes out comes back, where everything gets sampled and remixed, where everything has consequences, where social relationships co-evolve with material ones, where technology changes the very institutions producing it.[26]

The reflexivity of the world can be read positively (Giddens: "The reflexivity of modern social life consists in the fact that social practices are constantly examined and reformed in the light of incoming information about those very practices, thus constitutively altering their character"[27]) or less positively (Beck: "The gain in power from techno-economic 'progress' is being increasingly overshadowed by the production of risks"[28]). In any case, the process keeps getting more intensely reflexive. Once the marketplace takes the prime role in directing technological change, as it has since the 1970s, the process feeds back into itself.[29] A leading product of information technology is more information technology, all molded by the imperatives of a capitalist economy.

In terms of technological development, this is a virtuous circle—but it is still a circle, which accounts for the cultural paradox that the digital age engenders both a sense of liberating possibilities and a sense of oppression. Because of its socialized nature, information technology enhances social abilities, arguably our most human trait. But as information technology keeps reinforcing its dominance in terms defined by the market, other forms of sociability get selected out. A business-oriented society becomes more and more so as the system responds to its own feedback. In a world we have largely constructed, we keep encountering ourselves—but the projected "selves" we keep running into represent only a part of capacities, and arguably not the finest part.[30]

In the words of the historian Elting Morison, the historical challenge of our time is to build "a technological firmament that will really fit us" and "to organize a technological world we can live in."[31] History has not ended any more than nature is ended. Inevitably history will reveal itself to be much more powerful and unpredictable than the latest software upgrade. Inevitably nature will remind us of the fragility and impermanence of our control of the organic habitat. But history and nature are weakening as well-defined external frameworks that give meaning to human life. For most of human experience they were much more vast than the scale of individual existence, and their vastness provided solace in the case of nature

and hope in the case of history. When nature becomes so intermingled with humanity, and when historical change becomes defined as an extension of technological change, neither nature nor history provides a self-evident framework for human meaning and identity. The only larger framework left, it seems, is humanity itself.

NOTES

1. See Nicholas Negroponte, *Being Digital* (Knopf, 1995). (Negroponte, the longtime director of MIT's Media Lab, recently stepped down from that position.) The dichotomy of lifeworld and systems is associated with Jürgen Habermas; see Fred R. Dallmayr, "Life-World: Variations on a Theme," in *Life-World and Politics: Between Modernity and Post-modernity*, ed. S. White (University of Notre Dame Press, 1989), pp. 45–6 and notes. The dialectic of globalization and identity is a major theme of Manuel Castells, whose works are referenced throughout this book. The term "symbolic analyst" comes from Robert Reich, *The Work of Nations: Preparing Ourselves for 21st Century Capitalism* (Knopf, 1991). The quote about the disposable labor force is from Manuel Castells, *The Rise of the Network Society*, second edition (Blackwell, 2000 [1995]), p. 295.
2. This discussion is based on a book by Edward T. Layton Jr.: *The Revolt of the Engineers: Social Responsibility and the American Engineering Profession* (Press of Case Western Reserve University, 1971), p. 53. The historian of technology, Susan Douglas, reminds us that "the study of ideology is essential to the study of technology." She defines it as "the entire constellation of images, myths and ideas through which people come to understand the material world and their place in that world." Social groups are engaged in an ongoing struggle about the myths and meanings that will be accepted as the natural and obvious way to interpret social reality. "This," Douglas claims, "is a struggle invariably won by the dominant classes." (Susan J. Douglas, "Jürgen Habermas Meets Mel Kranzberg What Media Theory Has to Offer the History of Technology and Vice Versa," paper presented at annual meeting of Society of the History of Technology, 1987, p. 2.)
3. Karl Mannheim, *Ideology and Utopia: An Introduction to the Sociology of Knowledge* (Harcourt, Brace, 1936 [1929]), p. 203; see pp. 192–211
4. In listening for these words, their clustering, and their mutual defining, I am drawing on an honorable tradition in intellectual history—best exemplified by Raymond William's book *Culture and Society, 1780–1950* (Columbia University Press, 1958)—of analyzing the social and cultural implications of the industrial revolution. Williams highlighted the reflexive or circular relationship between key terms of social analysis in the nineteenth century (culture, society, industry, art, etc.) and the changes they supposedly described. In addition, Williams showed how key concepts tended to cluster, so that their meanings shifted in response to alterations in the others. See Marx, "Technology: The Emergence of a Hazardous Concept," *Social Research* 64 (1997), no. 3, p. 967.
5. Leo Marx, who has tracked the use of the word, calls "technology" a "hazardous concept," unstable and reflexive in meaning, and particularly dangerous when used as the subject of a sentence, implying that it is an independent, autonomous historical agent, as in "the archetypal sentence: 'Technology is changing the way we live.'" Marx argues that the word "technology" began to be used in the late nineteenth century to fill a "conceptual void" in explaining how the utilitarian, instrumentalist branch of human activity, directed at controlling the physical world, could so powerfully generate social change. (Marx, "Technology: The Emergence of a Hazardous Concept," pp. 968, 974–975, 977.) The naming of MIT in 1862 was one of the first public uses of the term. (The choice probably was influenced by Jacob Bigelow, a Boston botanist and physician, who had used the term in a Harvard lecture in 1828; he later became an MIT trustee.)
6. John Staudenmaier, "Perils of Progress Talk: Some Historical Considerations," in *Science, Technology, and Social Progress*, ed. S. Goldman (Lehigh University Press, 1989).
7. The process by which technological change began to displace historical change began in the seventeenth century with the controversy of the "ancients" and the "moderns." The "moderns" might concede equality or even superiority to the ancients in arts or philosophy, but they felt they clinched their case in arguing for progress when they emphasized that only the moderns possessed the compass, the printing press, gunpowder, and (most important) the experimental method.

were extrapolated to history as a whole; progress in what we now would call technology was cited as irrefutable evidence for historical progress from antiquity to the modern age. When technology entered the study of history, history itself began to be redefined as the study of the record of human progress. See Richard Foster Jones, *Ancients and Moderns: A Study of the Background of The Battle of the Books* (Washington University Studies, New Series in Language and Literature, no. 6, 1936). See also Hannah Arendt, *The Human Condition* (Double-day Anchor Books, 1959 [1958]), esp. pp. 119–133.

8. Anthony Giddens, *The Consequences of Modernity* (Stanford University Press, 1990), p. 51. Giddens puts "emptying out" in quotes in the text I am paraphrasing.
9. As Douglas reminds us in "Jürgen Habermas Meets Mel Kranzberg," the "'engineering of consent' is not a one-time event, in which meanings are clamped down and everyone accepts them." Consent is by no means automatic. It is an ongoing process in which people adjust their ideas and beliefs but also protest in a variety of ways. If the ideological framework becomes too much at odds with daily experience and with common sense, then the disparity has to be resolved for the social order to be stable: "The dominant ideology must incorporate certain criticisms and concerns while successfully marginalizing others, and illustrate that such concerns can be resolved best under the existing sociopolitical system."
10. Thomas Friedman makes this point when he uses a Lexus automobile and an olive tree to represent two sides of a technological divide: "Half the world seemed to be emerging from the Cold War intent on building a better Lexus, dedicated to modernizing, stream-lining and privatizing their economies in order to thrive in the system of globalization. And half of the world—sometimes half the same country, sometimes half the same person—was still caught up in the fight over who owns which olive tree." (*The Lexus and the Olive Tree*, Farrar, Straus and Giroux, 1999, p. 27.)
11. Alex Roland, email, March 5, 2000.
12. See Tom Standage, *The Victorian Internet: The Remarkable Story of the Telegraph and the Nineteenth Century's On-Line Pioneers* (Walker, 1998).
13. The concept of revolution migrated from physics (the path of the planets around the sun) to history (the overturning of British rule in America in the 1770s, the end of the monarchy in France in 1789), and then, in the early nineteenth century, in a deliberate analogy with events in France, was applied to machine production in the metaphor of the industrial revolution. When the metaphor of revolution is applied to industry, it implies a suddenness and completeness that might occur in the political arena (the king is executed, Robespierre rules) but does not apply in the arena of production. But to say that the metaphor is misleading is to miss the point. In the words of Leo Marx, "The whole issue becomes irrelevant once we recognize that we are dealing with a metaphor, and that its immense appeal rests, not on its capacity to describe the actual character of industrialization, but rather on its vivid suggestiveness. It evokes the uniqueness of the new way of life, as experienced, and, most important, it is a vivid expression of the affinity between technology and the great political revolution of modern times." (*The Machine in the Garden*, p. 187n.)
14. Steve Lohr, "The Future Came Faster in the Old Days," *New York Times*, October 5, 1997.
15. Lewis Mumford calls the invention of language "infinitely more complex and sophisticated ... than the Egyptian or Mesopotamian kit of tools" ("Technics and the Nature of Man," *Technology and Culture 7*, 1966, no. 3, p. 308).
16. On the term "habitat," see Gayle L. Ormiston, "Introduction," in *From Artifact to Habitat: Studies in the Critical Engagement of Technology*, ed. G. Ormiston (Lehigh University Press, 1990), p. 13.
17. Elting Morison, "Introductory Observations," in Morison, *Men, Machines, and Modern Times* (MIT Press, 1966), p. 16. See my discussion of the language and imagery of the artificial environment in *Notes on the Underground: An Essay on Technology, Society, and the Imagination* (MIT Press, 1990), pp. 1–8.
18. Manuel Castells (*The Rise of the Network Society*, p. 15) puts it into one long sentence: "Matter includes nature, human-modified nature, human-produced nature, and human nature itself, the labors of history forcing us to move away from the classic distinction between humankind and nature, since millenniums of human action have incorporated the natural environment into society, making us, materially and symbolically, an inseparable part of this environment."
19. Bill McKibben makes this point in *The End of Nature* (Random House, 1989).
20. Ulrich Beck, *Risk Society: Towards a New Modernity* (Sage, 1992 [1986]), p. 81.

21. Here I am paraphrasing Mumford ("Technics and the Nature of Man," p. 303), who writes: "In terms of the currently accepted picture of the relation of man to technics, our age is passing from the primeval state of man, marked by his invention of tools and weapons for the purpose of achieving mastery over the forces of nature, to a radically different condition, in which he will not only have conquered nature but detached himself completely from the organic habitat."
22. Mumford, "Technics and the Nature of Man," p. 303.
23. See Stanley H. Ambrose, "Paleolithic Technology and Human Evolution," *Science* 291 (2001), no. 5509, pp. 1748–1752: "Stone tool technology, robust australopithecines, and the genus *Homo* appeared almost simultaneously 2.5 [million years ago]" (p. 1748). The systematic use of fire evidently began between 1 million and 1.5 million years ago, and around 300,000 years ago "technological and cultural evolution accelerated" in a way that suggests "the emergence of true cultural traditions and cultural areas" (p. 1751). The stone tool technologies between 2.5 and 0.3 million years ago "are remarkable for their slow pace of progress … and for limited mobility and regional interaction … with the appearance of near-modern brain size, anatomy, and perhaps of grammatical language ~0.3 [million years ago], the pace quickens exponentially" (p. 1752).
24. See Peter F. Drucker, "The First Technological Revolution and Its Lessons," *Technology and Culture* 7 (1966), no. 2, pp. 143–151. This article is Drucker's presidential address to the Society for the History of Technology, presented on December 29, 1965, in San Francisco. He defines "the first technological revolution" as "the irrigation city, which then rapidly became the irrigation empire" (p. 143).
25. Mircea Eliade, *The Forge and the Crucible*, second edition (University of Chicago Press, 1978 [1956]), pp. 177–178. For a discussion of this passage in a different context, see Williams, *Notes on the Underground*, pp. 1–4.
26. The concept of reflexive modernization was articulated most clearly by Beck in *Risk Society*. See also Anthony Giddens's books *The Consequences of Modernity* (cited above) and *Modernity and Self-Identity in the Late Modern Age* (Stanford University Press, 1992).
27. Giddens, *The Consequences of Modernity*, p. 38.
28. Beck, *Risk Society*, pp. 12–13.
29. Under the governing structures of market relationships, the abiding imperative of capitalist accumulation—to increase profits—is accomplished in a few major ways. The circulation of capital can be maximized through increasing its extension in space (broadening the market) and its intensity over time (accelerating capital turnover). Also, production costs can be lowered, productivity can be raised, and new products can be brought to market. Particularly in the last three decades, all these ways of increasing profit are closely linked to technological innovation, which has become the primary driver of capitalist acquisition. See Castells, *The Rise of the Network Society*, pp. 68–69, 95. Castells comments, "Information processing is focused on improving the technology of information processing as a source of productivity, in a virtuous circle of interaction between the knowledge sources of technology and the application of technology to improve knowledge generation and information processing.... Or, more briefly, a prime characteristic of the information age is its circularity, characterized by the immediate application to its own development of technologies it generates…" (ibid., p. 32).
30. "Meanwhile man … exalts himself to the posture of lord of the earth. In this way the impression comes to prevail that everything man encounters exists only insofar as it is his construct. This illusion gives rise in turn to one final delusion: It seems as though man everywhere and always encounters only himself. Heisenberg has with complete correctness pointed out that the real must present itself to contemporary man in this way. *In truth, however, precisely nowhere does man today any longer encounter himself, i.e. his essence.*" (Martin Heidegger, "The Question Concerning Technology," in *The Question Concerning Technology and Other Essays*, Harper & Row, 1977, p. 27.) Hannah Arendt (*The Human Condition*, p. 237) quotes Heisenberg's observation about "man encountering only himself."
31. The first quote (from Morison, *From Know-How to Nowhere*) is cited by Thomas P. Hughes in "Memorials: Elting Morison, 1909–1995," *Technology and Culture 37* (1996), no. 4, p. 874. The second quote, also from "Memorials," is cited by Leo Marx (ibid., p. 867).

1.2 SOCIAL/POLITICAL PERSPECTIVES

1.2.1 I'D HAMMER OUT FREEDOM: TECHNOLOGY AS POLITICS AND CULTURE

RICHARD SCLOVE

Even though technologies are usually developed to meet a particular need or solve a specific problem, they often have secondary functions or unintended effects on our lives. In this selection, Richard Sclove is interested in how technologies exist within, are influenced by, and in turn exert influence on society. He claims that they often work in subtle ways to both define and set limits on human interaction.

The view of "technology as society" discussed in this selection states that although technology influences the social experience, it is not the sole determinant of how society develops. As a technology goes through various evolutionary stages, it assumes a greater role in our lives, and we become less aware of its influence and less likely to seek alternative ways of doing something. It is important, claims Sclove, for societies to look for alternative technologies more in keeping with humane ideals and aspirations.

FOCUS QUESTIONS

1. Identify one modern technology and discuss its development. In what ways has this technology changed the way you live from both a positive and negative perspective, and what future changes might occur that could have even more impact on your life?
2. Identify one environmental and organizational background imperative of a contemporary technology. How might those conditions have influenced that technology's development and its influence on our lives?
3. Sclove uses the hammer as an example of his concept of *polypotency*. What are some other examples of familiar artifacts that are polypotent? Discuss.

KEYWORDS

citizenship, democracy, primary function, social consequences, social structure

Technological innovations are similar to legislative acts or political foundings that establish a framework for public order that will endure over many generations.

LANGDON WINNER

What is technology? People ordinarily think of technology as machinery or gadgetry, as an economic factor of production, as know-how, as what engineers do, or as progress. Often they characterize technologies in terms of a single

Source: From *Democracy and Technology* by Richard E. Sclove. © 1995 The Guilford Press. Reprinted by permission.

intended function. What is a hammer? It's what someone uses to pound nails into boards. What is a telephone? It's a device that enables people to converse at a distance. Some technologies, however, have more than one intended function. Hammers, for example, can pound nails into boards but can also extract them. This is the core of the contemporary view of technology. People understand technologies in terms of a primary function—or, occasionally, several functions—that each is intended to accomplish.

Beyond this, our society has in the past few decades come to acknowledge that technologies tend to produce at least two general kinds of "secondary" or "unintended" effects. First, they generate environmental consequences: pollution, resource depletion, and ecosystem modification. Each of these may, in turn, have direct or indirect effects on human life. Second, they promote unintended social consequences—consequences that are generally mediated by economic markets (e.g., the replacement of workers by machines or the emergence of boomtowns). Thus common knowledge has it that technologies perform one or perhaps a few intended functions while also producing a limited range of unintended social and environmental consequences.

Although this view of technology is straightforward, it is also incomplete and misleading. It diverts attention from many significant aspects of technology, including some of central concern to democracy. By synthesizing recent technological criticism, the alternative view of technology introduced here incorporates the accepted view's sound insights but situates these within a broader perspective that recognizes technologies as a species of social structure.

The phrase "social structure" refers to the background features that help define or regulate patterns of human interaction. Familiar examples include laws, dominant political and economic institutions, and systems of cultural belief. Technologies qualify as social structures because they function politically and culturally in a manner comparable to these other, more commonly recognized kinds of social structures.

TECHNOLOGIES AS SOCIAL STRUCTURES

Ibieca, a Spanish village, found that its indoor plumbing came at the expense of community integration. That is an instance of a technology helping to structure social relations. Upsetting a traditional pattern of water use compromised important means through which the village had previously perpetuated itself as a self-conscious community. In the United States the automobile has played a somewhat similar role in disrupting prior patterns of community life.

These are not isolated cases; technologies designed for such mundane tasks as commuting to work or cooking food also routinely help constitute social systems of cooperation, isolation, or domination.

> Technology often embodies and expresses political value choices that, in their operations and effects, are binding on individuals and groups, whether such choices have been made in political forums or elsewhere.... Technological processes in contemporary society have become the equivalent of a form of law—that is, an authoritative or binding expression of social norms and values from which the individual or a group may have no immediate recourse.

Coercive Compliance

Technologies help regulate social behavior in part because they are themselves governed by both physical and political laws. For example, the operation of many technologies—such as automobiles, medical X-ray machines, or guns—is legally regulated. Thus their misuse can entail a socially enforced penalty.

However, whether or not they are governed by legal regulations, technologies generally embody a variety of other kinds of coercive mandates. The penalty for resisting these mandates may range from an informal reprimand ("Don't lick the food off your knife!") to economic loss or systemic failure (e.g., the gears in a conveyor belt jam, or a worker's hand is injured). These latter results are akin to the consequences befalling those who ignore physical laws (e.g., when someone literally walks on thin ice). Thus physical constraints, or accompanying legal and social sanctions, are

among the obvious means through which technologies help structure human behavior.

Subconscious Compliance

Sometimes technologies shape behavior and relationships less through brute compulsion than via subtle, psychological inducement. For example, social scientists have shown that the physical arrangement of chairs and tables strongly influences the kind of social interaction that occurs in schools, nursing homes, and hospitals. Yet the staff in those institutions had previously attributed behavior (including their own) entirely to the mix of personalities and psychological capabilities. They were surprised to learn that simply shifting the furniture could, for instance, help reanimate a seemingly moribund group of mentally impaired hospital inmates.

Opportunities and Constraints

Social structures are also ambiguous in that while they can restrict opportunities in some respects, they can—when appropriately designed—enhance them in others. For example, well-crafted laws help protect basic civil rights and, by providing a relatively stable and well-ordered social context, make it easier for people to realize their life plans. Besides creating novel opportunities and constraints, technologies also reconfigure prior patterns. For instance, within some offices and factories the proliferation of personal computer networks has enhanced lower level workers' chances to contribute to production decisions while simultaneously challenging midlevel managers' former domains of authority and autonomy. Once deployed, technologies can also aid or hinder the use of other technologies. For instance, telephone systems gradually displaced telegraph services but have more recently facilitated development of computer networks and long-distance data processing.

Background Conditions as Imperatives

In order to function, technologies require various environmental and organizational background conditions. A television set is only useful so long as viewers know how to operate it, it is protected from inclement weather, there is access to electricity, programs are being produced and distributed, and so on.

Frequently when individuals or groups acquire new technologies or technological facilities, they are at best only dimly conscious of the demands that effective operation will impose or require to be developed. Several years ago a town near mine in western Massachusetts approved construction of an industrial research center, hoping thereby to realize tax benefits. But no one asked beforehand the eventual costs (financial, environmental, and emotional) that the town would one day bear in order to accommodate both new research activities and the concomitant growth in commuting, ancillary employment, and residential population. These costs could include hazards associated with toxic waste disposal, future loss of open space to new housing, and the burden of upgrading roads, sewer lines, snowplowing capabilities, schools, and school bus lines.

To the extent that a given technology plays only a small part in one's life, maintaining the conditions needed for its operation may be of no particular concern. But as a person or society grows dependent on a technology, the necessary conditions of its operation loom as practical imperatives. The need to support these conditions represents a way in which technologies exert a profound structural social influence.

Technology as Structural for Nonusers

Often technologies exert comparably significant effects on people who neither operate nor use the technology in question. One clear example involves phenomena that economists label "spillover effects" or "externalities." Homeowners hear neighbors' radios, lawn mowers, or air conditioners; whole communities breathe noxious fumes from an industrial facility. Each person lives in an aesthetic landscape that reflects the aggregate technological choices made by other people or organizations. The psychological texture of our everyday life reflects the influence of

countless technological choices and practices in which we did not participate.

Moreover, often such spillover effects exert a structural influence that is dynamic and transformative. For instance, someone might choose not to purchase a power lawn mower to avoid its noise. However, after a few neighbors have bought theirs, this person may reconsider, thinking, "Since I'm suffering from the noise anyway, why not buy my own power mower and at least benefit from the convenience?" In this way each mower purchased contributes to a cycle that gradually transforms a neighborhood of quiet into one rent by the sound of churning engines.

Next reconsider the background conditions necessary for a technology to operate. Many of those conditions have a tremendous impact on lives even if individuals do not own the technology or use the technological service that establishes their *raison d'etre*. Suppose, to state the case dramatically, that as a citizen of a modern nation a woman opts for a relatively self-sufficient mode of life: She refuses to own a car, uses solar collectors on the roof of her home, and plants a large vegetable garden in her yard. What has she accomplished? Something, certainly, but the texture of her world still reflects the existence not only of cars and their immediate culture, but of roadways, automobile manufacturing and marketing systems, oil refineries, electric generating facilities, agribusiness, the private or public bureaucracies that manage these things, and their often tumultuous politics. That is part of what it means to say that technologies are social structures. The aggregate result of a society's many technological choices in one way or another affects every member.

Communicative and Cultural Systems

Apart from materially influencing social experience, technologies also exert symbolic and other cultural influences. This is true not only of technologies explicitly called communications devices (e.g., cellular phones, televisions, and radios), but of all technologies.

For example, modern sofas generally have two compelling technical or economic rationale for this design (an affordable, seamless sofa is an easily conceived alternative—as seamless mattresses and Japanese futon sofa-beds attest). Rather, separate sofa cushions define distinct personal spaces and thus respect—but also help to perpetuate—modern Western culture's emphasis on individuality and privacy.

Technologies even play transformative roles within psychological development. For example, earlier this century Swiss psychologist Jean Piaget determined that young children distinguish living from nonliving things according to whether or not the things move, and—as the children develop psychologically—then according to whether things move by themselves or are moved by an outside force. However, more recently, social psychologist Sherry Turkle found that children who play with computer toys that appear to "talk" and "think" develop different criteria for distinguishing "alive" from "not alive." Instead of relying on physical criteria (such as motion), they invent psychological criteria and hypotheses ("Computers are alive because they cheat" or "Computers are not alive because they don't have feelings"). Children's developmental trajectories, including their conceptions of self and moral reasoning, are transformed as a result of their interactions with these machines.

The process that Turkle described with respect to computer toys is a specific instance of a much more general phenomenon. As they reconfigure opportunities and constraints for action, and function simultaneously as symbols and latent communicative media, technologies also reconfigure opportunities and constraints for psychological development.

Macropolitics: Technology *and* Society Versus Technology *as* Society

Many scholars have described cases in which technologies exert a macrolevel influence on societies. Consider historians who focus on the social role of just one or two important technologies at a time. Large scale dams and irrigation systems may have played a decisive role in the creation and mainte-

now-famous story of the role of the stirrup in the development of European feudalism: stirrups made possible mounted shock combat, which led in turn to heavy full-body armor, heraldry, chivalry, stronger horse breeds, more efficient plowing methods, and so forth. In America, railroads helped establish national markets; promoted coal mining, steelmaking, and the widespread adoption of steam power; provided an influential model of geographically dispersed, hierarchically managed corporate organization; contributed to the adoption of standardized timekeeping; and served as a dominant metaphor with which Americans interpreted their entire civilization.

More recently in the United States one role of new technologies has been to provide grounds for the growth of the federal government, through the proliferation of such agencies as the Federal Communications Commission for regulating telecommunications, the Federal Aviation Administration for regulating the airline industry, the Nuclear Regulatory Commission and the Department of Energy for administering aspects of national energy production and nuclear weapons development, and the like.

In each of these instances, technological innovation plays a role in establishing, transforming, or maintaining states or societies at the macrolevel. Langdon Winner has explored the further hypothesis that the entire ensemble of modern technological systems—including the background conditions required to keep them operating—tends to promote centrally coordinated, technocratic social administration.

Hence there are numerous examples in which technologies affect societies or states in ways that have macrostructural implications. However, this formulation—while both true and dramatic—nonetheless misses the force of this chapter's earlier analysis. Technologies function politically and culturally as social structures by coercing physical compliance; prompting subconscious compliance; constituting systems of social relations; establishing opportunities and constraints for action and self-realization; promoting the evolution of background conditions; affecting nonusers; shaping communication, psychological development, and culture; generally and constituting much of the world within which lives unfold.

Considering all of the preceding functions and effects together, it would be fairer to say that technologies do not merely *affect* society or states, they also *constitute* a substantial portion of societies and states. That, too, is part of what it means to be a social structure. Recognizing the many respects in which technologies contribute to defining who people are, what they can and cannot do, and how they understand themselves and their world should dispel the common myth that technologies are morally or politically neutral.

Influential, Not Determining

Technologies "structure" social relations in that they shape or help constitute—but do not fully determine—social experience. Water pipes and washing machines did not, for example, literally force Ibiecans to stop gathering at their village's central fountain and washbasin, but instead altered the system of inducements and interdependencies that formerly made such gathering occur naturally.

Aside from the possibility of rejecting or retiring a particular technology, there is always a margin of flexibility in how existing technological artifacts may be used or operated, or in what activities may occur in conjunction with them. This margin is finite, and its extent varies from one technology to the next and over time, but it nevertheless exists. For example, while a conventional assembly line provides only highly restricted opportunities to vary work routines at each station, it does not materially prevent workers from rotating jobs among work stations.

Context-Dependency

Developing a railroad network helped catapult the United States to global economic preeminence, but Britain developed railroads earlier and yet nonetheless gradually lost its world economic predominance. Thus railroads (or other technologies) are socially consequential, but how and why they matter depends on the precise technologies in

Moreover, just as social context—including, among other things, a society's preexisting technological order—regulates each technology's material functions and effects, it also regulates a technology's communicative functions and cultural meanings. A few decades ago a belching smokestack symbolized progress. Today—in a different historical context—the same smokestack is more likely to evoke distress or even outrage.

Finally, one important influence on a technology's functions and effects is the minds and culture of people. Nineteenth-century high-wheeler bicycles were perceived by athletic young men as virile, high-speed devices. But to some women and elderly men the same devices signified personal danger. Indeed, conflicting perceptions of the high-wheeler proved consequential to its subsequent technological development. Its perception as a "macho machine" prompted new bicycle designs with ever higher front wheels. The competing perception of the high-wheeler as an "unsafe machine" prompted designs with smaller front wheels, different seat placement, or higher rear wheels. Thus to understand the social function, meaning, and evolution of the high-wheeler, it is essential to explore its psychological and cultural context.

Public controversies concerning technology offer another occasion for observing the role of culture and cognition in establishing a technology's context, and hence its social role. For example, during the 1970s nuclear engineers and electric utility executives generally viewed centralized production of electricity as a critical social need and essential to the concept of commercial nuclear power. To them an alternative to nuclear power needed to be another means of performing this critical function. But other energy policy analysts saw the expanded production of electricity as so inessential that a perfectly viable alternative could be a panel of foam wall insulation that did not generate any electricity.

In evidence here are fragments of a social process of contesting or negotiating what is or is not to count as an essential function of a technology and hence as an alternative. Thus, when technological consequences or meanings become controversial, processes through which technologies are culturally constituted may emerge openly.

Contingent Social Products

There is residual variability in the structural effects associated with any deployed technology—*within a particular social context and even more so among different contexts*. However, a technology's greatest flexibility exists before its final deployment, when artifacts and their accompanying social organization are being conceived and designed.

Technologies do not just appear or happen; they are contingent social products. Thus it is possible, both before and after the fact, to imagine alternative designs. The process by which one set of designs rather than another comes to fruition is influenced by prevailing social structures and forces, including the preexisting technological order. However, this process also reflects explicit or tacit social choices, including political negotiations or struggles.

For example, it is hard to imagine a modern home without an electric refrigerator, but had the accidents of competing corporate resources played out slightly differently, gas-powered refrigerators that would have run more reliably and quietly could have been the norm. Other feasible alternatives in household technology harbored the potential for even more dramatic social effects. Moreover, although today people think of the guiding impulses behind technological development as necessarily being profit, convenience, or military advantage, throughout history religious or aesthetic motivations have often been just as significant.

Thus there are many potential, competing technological pathways, and each is socially developed. But the flexibility associated with a given technology, or with other social structures, tends to diminish with time. After a society has habituated itself to one technology, alternatives tend to become less accessible. Once designed and deployed, a technology, like a law or a political institution, tends—if it is going to endure—gradually to become integrated into larger systems of functionally interdependent artifacts and organizations and then to influence the design of

RICHARD SCHOVE | Pd Hammer Out Freedom: Technology as Politics and Culture 75

subsequent technologies, laws, and institutions such that the latter all tend to depend on the continued existence of the former. Thus, owing to the accompanying evolution of supporting custom, entrenched interest, and various sunk costs, it is often difficult to achieve radical design alterations once an initial decision has been implemented. A further factor reducing the flexibility of technologies is that they exhibit some of the pure physical recalcitrance that comes with material embodiment. Hence, both technologies and other social structures, once they have come into existence, tend to endure. However, technologies exhibit a remaining characteristic that tends to distinguish them from other social structures and to increase their relative political salience: *polypotency*.

POLYPOTENCY

Technologies function as social structures, but often independently of their (nominally) intended purposes. This is one of the phenomena that the conventional view of technology obscures. The same obfuscation is reflected in studies that profess a broad interest in the political effects of technology but that discuss only technologies designed explicitly to function politically (such as telecommunications, military and police technologies, voting machines, or computer databases). Such technologies indeed function politically, but everyone knows that. That is these technologies' announced purpose. Harder to grasp is the truth that all technologies are associated with manifold latent social effects and meanings, and that it is largely in virtue of these that technologies come to function as social structures. In other words, technologies exhibit superfluous efficacy or "polypotency" in their functions, effects, and meanings. (The word *polypotency*, meaning "potent in many ways," is introduced here for want of a better existing term. The unfamiliarity wears off quickly if one contrasts it with omnipotence, meaning, literally, "potent in all ways.")

For example, when a man uses an ordinary hammer to pound nails, he also learns about the texture and structural properties of materials, he exercises and develops his muscles, he improves his hand-eye coordination, and he generates noise, all while stressing and wearing the hammer itself. As his competence at hammering grows, he feels his self-respect affirmed and approved. At another level, his activity resonates with half-conscious memories of primeval myths about Vulcan and Thor. He is also reminded of the blacksmith and the mythology of the American frontier. He thinks of a judge's gavel, the hammer as a symbol of justice, and a song popularized by the folk singing trio Peter, Paul, and Mary.

Where did the hammer come from? Somebody chopped down a tree and fashioned the handle. Others located and extracted iron ore. Some of that ore was refashioned into a hammer head. If a man touches his tongue to the hammer, with the taste of oxidized iron he senses fleetingly a former age when once-independent craftsmen and farmers first found themselves working under strict supervision in a factory. When he was a child, an uncle first taught him to use a hammer. Now when he hefts a hammer, he feels embedded in a historical relationship with this and other hammers and with the development of the concept of hammers and technology in general.

The hammer's immediate social context of use can vary. The man may work alone, on a project with others, or in a room where each person pursues a different project. He may or may not choose his task; he may or may not earn a wage. Depending on the precise social context of its use, the hammer means different things to him, he sees it differently, and it helps disclose the world to him in different ways. Likewise, his style of using the hammer discloses to others much about his character, competence, and mood.

The hammer differs from a partially automated assembly line in that the latter requires and helps coordinate the simultaneous efforts of many workers. But a hammer also establishes certain limiting possibilities on the social conditions of its use. Hammers have only one handle. They are not designed to permit the type of close collaboration that is possible through computer networks or necessary when using a long, two-handled saw.

The material result of the man's activity is likely to include some bent nails, scrap wood, a hearty appetite, maybe a bruised thumb, a few

sore but marginally strengthened muscles, some excess exhalation of carbon dioxide, perspiration, and a product that becomes part of the humanly shaped world. So, is the nail entering the board necessarily the most important feature of the activity called "hammering"? Hammers, like all technologies, are polypotent in their social functions, effects, and meanings.

Today's accepted view of technology takes a step toward acknowledging polypotency by speaking of technologies' unintended or secondary consequences. However, the term "polypotency" is helpful in not presuming that one knows automatically which of a technology's many functions or meanings are the most important or even which are intended. Many social historians of technology have, for example, argued that a latent but intended function of some innovations in manufacturing technology has been to substitute low-paid unskilled workers for higher-paid skilled workers, discipline the remaining workforce, and weaken unions.

It is furthermore useful to introduce the term "focal function" to refer to a technology's (ostensibly) intended purpose. "Nonfocal" then denotes its accompanying complex of additional—but often recessive—functions, effects, and meanings. Thus, nineteenth-century New England schoolhouses' focal function was to provide a space for educational instruction, whereas one of their nonfocal functions was to help generate—in part via the symbolism of churchlike architectures—a relatively docile workforce.

Occasionally technologies function as social structures precisely by virtue of their focal purpose. For instance, weapons function coercively because they are designed to do just that. But more often and more subtly, it is technologies' latent polypotency that accounts for their structural performance. This is illustrated by many previous examples, ranging from sofa cushions (which help to latently reproduce our culture's sense of privacy) to computer toys (which unexpectedly alter children's psychological development). Even technologies focally designed to function structurally are apt to structure nonfocally as well.

For instance, nuclear weapons are designed focally to coerce, deter, or destroy other societies, but they contribute nonfocally to legitimating authoritarian government institutions within the societies that possess them. Marshall McLuhan popularized this truth as it applies specifically to technologies focally designated as communications devices: "The medium *is* the message." In other words, the technical means of focally delivering a message can, owing to polypotency, matter more than the message itself.

Moreover, often groups of focally unrelated technologies interact latently to produce a structural effect that no one of them could accomplish alone. Distinct sofa cushions would not help establish cultural norms of privacy and individualism were they not part of a complex of artifacts and ritual behavior that contribute jointly toward that same result. (Other artifacts in the complex with sofa cushions include individual eating utensils, private bedrooms, telephone receivers designed to accommodate one person at a time, and so forth.) In short, to achieve social insight and efficacy, it is essential to consider all the different artifacts and practices that comprise a society's technological order.

There are important functional equivalences between technologies and nontechnological social structures (e.g., legal statutes, government agencies, and large corporations). All represent enduring social products that shape subsequent social experience. However, there are also differences, revolving around contrasting levels of social understanding with respect to each.

First, laws and political and economic institutions are contingent social products, and at some level everyone knows this truth. Societies evolve these things through formal political or juridical processes, and it is commonly understood that alternative choices are possible. In contrast, people are prone to misperceive a society's technologies as inevitable, that is, as naturally determined rather than socially shaped and chosen.

Second, laws and other formally evolved social structures are commonly understood to function as social structures. That is their explicit purpose. Certainly, they can also be implicated in

the production of various unintended social consequences. Prohibition-era laws were enacted to stop alcohol production, not to drive it underground and contribute to the expansion of organized crime. However, people at least expect that legal statutes and institutions will—because that is their intent—in some way shape social interaction and history. In contrast, people ordinarily expect most technologies to prove structurally inconsequential, and—because focally most of them do—this expectation appears confirmed. But here is where appearances deceive, insofar as it is frequently a technology's nonfocal aspects alone that conspire to manifest profound structural consequences.

Hence, although technologies are as consequential as other social structures, people tend to be more blind both to the social origins of technologies and to their social effects. This dual blindness is partly due to certain myths or misconceptions, such as the myth that technologies are autonomous self-contained phenomena and the myth that they are morally neutral. It is also inculcated through modern technologies themselves, via both their style and their social process of design.

These dual misperceptions concerning technologies actually enhance their relative structural significance, because they enable technologies to exert their influence with only limited social awareness of how, or even that, they are doing so. This helps explain why people are prone to resign themselves to social circumstances established through technological artifice and practices that they might well reject if the same results were proposed through a formal political process.

So long as their social origin, effects, and dynamics remain so badly misperceived, technologies will not suffer the same liability as would, say, functionally comparable laws or economic institutions, of being challenged on the grounds that they are politically or culturally unacceptable. Furthermore, societies will fail to develop the capacity to seek other technologies more consonant—both focally and nonfocally—with their members' ideals and aspirations.

1.2.2 DEMOCRATIC RATIONALIZATION

ANDREW FEENBERG

In his "Ode, inscribed to William H. Channing" the American transcendentalist poet Ralph Waldo Emerson wrote:

> "Things are in the saddle and ride mankind.
> There are two laws discrete,
> Not reconciled,
> Law for man, and law for thing;
> The last builds town and fleet,
> But it runs wild,
> And doth the man unking."

These verses capture the sense that many of us have of being controlled by and at the mercy of technology rather than controlling it. The question is: Does technology determine social conditions and practices? Or does society shape technology?

Source: Feenberg, Andrew, Afterword by Michel Callon, Foreword by Brian Wynne., *Between Reason and Experience: Essays in Technology and Modernity*, pp. 5–23, 26–29, © 2010 Massachusetts Institute of

In this selection, a leading philosopher of technology, Andrew Feenberg, Professor and Canada Research Chair in Philosophy of Technology at Simon Fraser University, contends against the idea of technological determinism. Drawing on historical examples of the social construction of technology [SCOT], he argues social and ethical values have shaped the path of technological development. Sociotechnical practices are flexible and can select among different technical means of achieving a particular purpose. But society can also provide cultural meanings for particular technologies or artifacts that make a difference in the things themselves. In the end, society is really in the saddle, even though any particular individual may not be able to determine which technologies become adopted and widely diffused.

Feenberg's constructivist perspective has potential radical implications for social theory generally; it challenges the old distinction between capitalism and socialism by suggesting that the important question really is whether the power to shape technology resides only within business and technical elites, or if the design and control of technologies are effectively democratized. Technology, on this view, is an arena of social and political struggle in which the values and interests of different individuals and groups in society contend for control over the shape and characteristics of the built environment.

FOCUS QUESTIONS

1. Discuss how those supporting and opposing the Factory Bill in 1844 in England framed their arguments and the role of determinism in those debates.
2. What were some factors that affected the design of the bicycle and how do the views of engineers and managers differ from those "studying the social role of technical objects and the lifestyles they make possible?"
3. What contemporary device posed a modern dilemma similar to that created by the bicycle, how did the French telephone company attempt to "sell" this new technology to the masses, and to what degree were they successful?
4. Discuss the ways in which cultural norms impact technological development and the role of rationalization as our modern horizon.

KEYWORDS

constructivism, cultural norms, determinism, dystopian modernity, extrapolitical domain, hegemony, political democracy, public power, rationalization, technique, technological nondeterministic theory, workplace democracy

THE LIMITS OF DEMOCRATIC THEORY

Technology is one of the major sources of public power in modern societies. So far as decisions affecting our daily lives are concerned, political democracy is largely overshadowed by the enormous power wielded by the masters of technical systems: corporate and military leaders and professional associations of groups such as physicians and engineers. They have far more to do with control over patterns of urban growth, the design of dwellings and transportation systems, the selection of innovations, and our experience as employees, patients, and consumers than do all the governmental institutions of our society put together.

Marx saw this situation coming in the middle of the nineteenth century. He argued that traditional democratic theory erred in treating the economy as an extrapolitical domain ruled by natural laws such as the law of supply and demand. He claimed that we will remain disenfranchised and alienated so long as we have no say in industrial decision making. Democracy must be extended from the political domain into the world of work. This is the underlying demand behind the idea of socialism.

Modern societies have been challenged by this demand for over a century. Democratic political theory offers no persuasive reason of principle to reject it. Indeed, many democratic theorists endorse it (Cunningham 1987). What is more, in a number of countries, socialist parliamentary victories or revolutions have brought to power parties dedicated to achieving it. Yet today we do not appear to be much closer to democratizing industrialism than in Marx's time.

This state of affairs is usually explained in one of the following two ways.

Technology is determining On the one hand, common sense argues that modern technology is incompatible with workplace democracy. Democratic theory cannot reasonably press for reforms that would destroy the economic foundations of society. For evidence, consider the Soviet case: although they were socialists, Lenin and his successors did not democratize industry, and even at its most liberal, the democratization of Soviet society extended only to the factory gate. Today in the ex-Soviet Union, everyone still agrees on the need for authoritarian industrial management.

Technology is neutral On the other hand, a minority of radical theorists claim that technology is not responsible for the concentration of industrial power. That is a political matter, due to the victory of capitalist and communist elites in struggles with the underlying population. No doubt modern technology lends itself to authoritarian administration, but in a different social context it could just as well be operated democratically.

In what follows, I will argue for a qualified version of this second position, somewhat different from both the usual Marxist and radical democratic formulations. The qualification concerns the role of technology, which I see as *neither* determining nor as neutral. I will argue that modern forms of hegemony are based on a specific type of technical mediation of a variety of social activities, whether it be production or medicine, education or the military, and that, consequently, democratization requires radical technical as well as political change.

This is a controversial position. Political theorists usually limit the proper application of the concept of democracy to the state. By contrast, I believe that unless democracy can be extended beyond its traditional bounds into the technically mediated domains of social life, its use value will continue to decline, participation will wither, and the institutions we identify with a free society will gradually disappear.

Let me turn now to the background of my argument. I will begin by presenting an overview of various theories that claim that technologically advanced societies require authoritarian hierarchy. These theories presuppose a form of technological determinism that is refuted by historical and sociological arguments I will briefly summarize. I will then present a sketch of a nondeterministic theory of modern society I call "critical theory of technology." This alternative approach emphasizes the impact of contextual aspects of technology on design ignored by the dominant view. I will argue that technology is not just the rational control of nature; both its development and impact are intrinsically social. I will then show that this view undermines the customary reliance on efficiency as an explanation of technological development in both optimistic and dystopian accounts of modernity. This conclusion, in turn, opens broad possibilities of change foreclosed by the usual understanding of technology.

Dystopian Modernity

Max Weber's famous theory of rationalization is the original argument against industrial democracy. The title of this chapter implies a provocative reversal of Weber's conclusions. He defined rationalization as the increasing role of calculation and control in social life, a trend leading to what he called the "iron cage" of bureaucracy (Weber 1958, 181–182). "Democratic" rationalization is thus a contradiction in terms.

Once traditionalist struggle against rationalization has been defeated, further resistance in a Weberian universe can only affirm an irrational life force against routine and drab predictability. This is not a democratic program but a romantic anti-dystopian one, the sort of thing that is already foreshadowed in Dostoyevsky's *Notes from Underground* and various back-to-nature ideologies.

My title is meant to reject the dichotomy between rational hierarchy and irrational protest implicit in Weber's position. If authoritarian social hierarchy is truly a contingent dimension of technical progress, as I believe, and not a technical necessity, then there must be an alternative rationalization of society that democratizes rather than centralizes control. We need not go underground or native to preserve threatened values such as freedom and individuality.

But the most powerful critiques of modern technological society follow directly in Weber's footsteps in rejecting this possibility. I am thinking of Heidegger's formulation of "the question of technology" and Ellul's theory of "the technical phenomenon" (Heidegger 1977; Ellul 1964). According to these theories, we have become little more than objects of technique, incorporated into the mechanism we have created. The only hope is a vaguely evoked spiritual renewal that is too abstract to inform a new technical practice.

... But first, to advance my own argument, I will concentrate on the principal flaw of dystopianism, the identification of technology in general with the specific technologies that have developed in the last two centuries in the West. These are technologies of conquest that pretend to an unprecedented autonomy; their social sources and impacts are hidden. I will argue that this type of technology is a particular feature of our society and not a universal dimension of modernity as such.

Technological Determinism

Determinism rests on the assumption that technologies have an autonomous functional logic that can be explained without reference to society. Technology is presumably social only through the purpose it serves, and purposes are in the mind of the beholder. Technology would thus resemble science and mathematics by its intrinsic independence of the social world.

Yet unlike science and mathematics, technology has immediate and powerful social impacts. It would seem that society's fate is at least partially dependent on a nonsocial factor that influences it without suffering a reciprocal impact. This is what is meant by "technological determinism." A deterministic view of technology is commonplace in business and government, where it is often assumed that technical progress is an exogenous force influencing society rather than an expression of changes in culture and values.

Dystopian visions of modernity are also deterministic. If we want to affirm the democratic potentialities of modern industrialism, we will therefore have to challenge their deterministic premises, the thesis of unilinear progress, and the thesis of determination by the base.

1. According to determinism, technical progress follows a unilinear course, a fixed track, from less to more advanced configurations. Although this seems obvious from a backward glance at the development of any familiar technical object, in fact it is based on two claims of unequal plausibility: first, that technical progress proceeds from lower to higher levels of development; and second, that that development follows a single sequence of necessary stages. As we will see, the

first claim is independent of the second and not necessarily deterministic.

2. Determinism also affirms that social institutions must adapt to the "imperatives" of the technological base. This view, which no doubt has its source in a certain reading of Marx, is now part of the common sense of the social sciences (Miller 1984, 188–195). I will discuss one of its implications in detail: the supposed "trade-off" between prosperity and environmental values.

These two theses of technological determinism present decontextualized, self-generating technology as the foundation of modern society. Determinism thus implies that our technology and its corresponding institutional structures are universal, indeed, planetary in scope. There may be many forms of tribal society, many feudalisms, even many forms of early capitalism, but there is only one modernity, and it is exemplified in our society for good or ill. Developing societies should take note: as Marx once said, calling the attention of his backward German compatriots to British advances: "*De te fabula narratur*"—of you the tale is told (Marx 1906 reprint, 13).

CONSTRUCTIVISM

The implications of determinism appear so obvious that it is surprising to discover that neither of its two theses withstands close scrutiny. Yet contemporary sociology undermines the thesis of unilinear progress, while historical precedents are unkind to the thesis of determination by the base.

Recent constructivist sociology of technology grows out of social studies of science (Bloor 1991, 175–179; Latour 1987). I employ the term "constructivism" loosely to refer to the theory of large-scale technical systems, social constructivism, and actor-network theory. They have in common an emphasis on the social contingency of technical development. They challenge the traditional view of the autonomy of technology and study it much as one might an institution or a law. The specifics of these methodologies are not relevant here, but this general approach lends support to the critical theory of technology.

Constructivism challenges our tendency to exempt scientific theories from the sort of sociological examination to which we submit nonscientific beliefs. It affirms the "principle of symmetry," according to which all contending beliefs are subject to the same type of social explanation regardless of their truth or falsity. A similar approach to technology rejects the usual assumption that technologies succeed on purely functional grounds.

Constructivism argues that theories and technologies are underdetermined by scientific and technical criteria. Concretely, this means two things: first, there is generally a surplus of workable solutions to any given problem, with social actors making the final choice among several viable options; and second, the problem definition often changes in the course of solution.

Trevor Pinch and Wiebe Bijker illustrate these points with the example of the bicycle. In the late nineteenth century, before the present form of the bicycle was fixed, design was pulled in several different directions. Some customers perceived bicycling as a competitive sport, while others had an essentially utilitarian interest in transportation. Designs corresponding to the first definition had high front wheels that were rejected as unsafe by the second type of rider. They preferred the "safety" with two equal-sized low wheels. With the introduction of inflatable tires the low wheelers won out, and the entire later history of the bicycle down to the present day stems from that line of technical development. Technology is not determining in this example; on the contrary, the "different interpretations by social groups of the content of artifacts lead via different chains of problems and solutions to different further developments" (Pinch and Bijker 1989, 42).

Pinch and Bijker call this variability of goals the "interpretative flexibility" of technologies. What a technology *is* depends on what it is *for*, and that is often in dispute. The flexibility of technologies is greatest at the outset and diminishes as the competition between alternatives is sorted out. Finally, closure is achieved in the consolidation of a standard design capable of prevailing for an

extended period. This is what happened to the bicycle, the automobile, and most of the familiar technologies that surround us.

In the case of the bicycle, the "safety" design won out, and it benefited from all the later advances. In retrospect, it seems as though the high wheelers were a clumsy and less efficient stage in a progressive development leading through the old "safety" bicycle to current designs. In fact the high wheeler and the "safety" shared the field for years, and neither was a stage in the other's development. The high wheeler represents a possible alternative path of bicycle development that addressed different problems at the origin. The defeated alternative was left frozen in time like a dinosaur fossil and so appears obviously inferior today in a typical illusion of progress.

Determinism is a species of Whig history that tells the story as though the end was inevitable by projecting the abstract technical logic of the finished object back into the past as the *telos* of development. That approach confuses our understanding of the past and stifles the imagination of a different future. Constructivism can open up that future, although its practitioners have hesitated so far to engage the larger social issues implied in their method.

INDETERMINISM

If the thesis of unilinear progress falls, the collapse of the notion of determination by the technological base cannot be far behind. Yet it is still frequently invoked in contemporary political debates.

I shall return to these debates later in this chapter. For now, let us consider the remarkable anticipation of current attitudes in the struggle over the length of the workday and child labor in mid-nineteenth-century England. The debate on the Factory Bill of 1844 was entirely structured around the opposition of technological imperatives and ideology. Lord Ashley, the chief advocate of regulation, protested that "The tendency of the various improvements in machinery is to supersede the employment of adult males, and substitute in its place, the labor of children and females. What will be the effect on future generations, if their tender frames be subjected, without limitation or control, to such destructive agencies?"

He went on to deplore the decline of the family consequent upon the employment of women, which "disturbs the order of nature" and deprives children of proper upbringing. "It matters not whether it be prince or peasant, all that is best, all that is lasting in the character of a man, he has learnt at his mother's knees." Lord Ashley was outraged to find that "females not only perform the labor, but occupy the places of men; they are forming various clubs and associations, and gradually acquiring all those privileges which are held to be the proper portion of the male sex … they meet together to drink, sing, and smoke; they use, it is stated, the lowest, most brutal, and most disgusting language imaginable …"

Proposals to abolish child labor met with consternation on the part of factory owners, who regarded the little worker as an "imperative" of the technologies created to employ him. They denounced the "inefficiency" of using full-grown workers to accomplish tasks done as well or better by children, and they predicted all the usual catastrophic economic consequences—increased poverty, unemployment, loss of international competitiveness—from the substitution of more costly adult labor. Their eloquent representative, Sir J. Graham, therefore urged caution: "We have arrived at a state of society when without commerce and manufactures this great community cannot be maintained. Let us, as far as we can, mitigate the evils arising out of this highly artificial state of society; but let us take care to adopt no step that may be fatal to commerce and manufactures."

He further explained that a reduction in the workday for women and children would conflict with the depreciation cycle of machinery and lead to lower wages and trade problems. He concluded that "in the close race of competition which our manufacturers are now running with foreign competitors … such a step would be fatal…." Regulation, he and his fellows maintained in words that echo still, is based on a

"false principle of humanity, which in the end is certain to defeat itself." One might almost believe that Ludd had risen again in the person of Lord Ashley: the issue is not really the regulation of work, "but it is in principle an argument to get rid of the whole system of factory labor." Similar protestations are heard today on behalf of industries threatened with what they call environmental "Luddism."

Yet what actually happened once the regulators imposed limitations on the workday and expelled children from the factory? Did the violated imperatives of technology come back to haunt them? Not at all. Regulation led to an intensification of factory labor that was incompatible with the earlier conditions in any case. Children ceased to be workers and were redefined socially as learners and consumers. Consequently, they entered the labor market with higher levels of skill and discipline that were soon presupposed by technological design. A vast historical process unfolded, partly stimulated by the ideological debate over how children should be raised and partly economic. It led eventually to the current situation in which nobody dreams of returning to cheap child labor in order to cut costs, at least not in the developed countries.

This example shows the tremendous flexibility of the technical system. It is not rigidly constraining but on the contrary can adapt to a variety of social demands. This conclusion should not be surprising, given the responsiveness of technology to social redefinition discussed previously. In sum, technology is just another dependent social variable, albeit an increasingly important one, and not the key to the riddle of history.

Determinism, I have argued, is characterized by the principles of unilinear progress and determination by the base; if determinism is wrong, then research must be guided by two contrary principles. In the first place, technological development is not unilinear but branches in many directions and could reach generally higher levels along several different tracks. In the second place, technological development is not determining for society but is overdetermined by both technical and social factors.

The political significance of this position should also be clear by now. In a society where determinism stands guard on the frontiers of democracy, indeterminism "enlarges the field of the possible." If technology has many unexplored potentialities, no technological imperatives dictate the current social hierarchy. Rather, technology is a scene of social struggle, a "parliament of things," on which civilizational alternatives contend (Latour 1993).

INTERPRETING TECHNOLOGY

In the next sections of this chapter, I would like to present several major themes of a nondeterminist approach to technology. The picture sketched so far implies a significant change in definition. Technology can no longer be considered as a collection of devices nor, more generally, as the sum of rational means. These are tendentious definitions that beg the question of technology's social significance and involvements.

Insofar as it is social, technology ought to be subject to interpretation like any other cultural artifact, but it is generally excluded from humanistic study. We are assured that its essence lies in a technically explainable function rather than a hermeneutically interpretable meaning. At most, humanistic methods might illuminate extrinsic aspects of technology, such as packaging and advertising, or popular reactions to controversial innovations such as nuclear power. Technological determinism draws its force from this attitude. If one ignores most of the connections between technology and society, it is no wonder that technology then appears to be self-generating.

Technical objects have two hermeneutic dimensions that I call their *social meaning* and their *cultural horizon*. The role of social meaning is clear in the bicycle case. We have seen that the design of the bicycle was decided by a contest of interpretations: Was it to be a sportsman's toy or a means of transportation? Design features such as wheel size signified it as one or another type of object while also suiting it to its function.

It might be objected that this is merely an initial disagreement over functions with no hermeneutic significance. Once the object is stabilized, the engineer has the last word on its nature, and the humanist interpreter is out of luck. This is the view of most engineers and managers; they readily grasp the concept of "function," but they have no use for "meaning."

In fact the dichotomy of function and meaning is a product of modern technical cultures, which are themselves rooted in the structure of the modern economy. The concept of "function" strips technology bare of social contexts, focusing engineers and managers on just what they need to know to do their job. A fuller picture is conveyed, however, by studying the social role of technical objects and the lifestyles they make possible. That picture places the abstract notion of "function" in its concrete social context. It makes technology's contextual causes and consequences visible rather than obscuring them behind an impoverished functionalism.

The functionalist point of view yields a decontextualized temporal cross-section in the life of the object. As we have seen, determinism claims implausibly to be able to get from one such momentary configuration of the object to the next on purely technical terms. But in the real world all sorts of unpredictable attitudes crystallize around technical objects and influence later design changes. The engineer may think these are extrinsic to the device he or she is working on, but they are its very substance as a historically evolving phenomenon.

These facts are recognized to a certain extent in the technical fields themselves. With computers, we have a contemporary version of the dilemma of the bicycle discussed earlier. Progress of a generalized sort in speed, power, and memory goes on apace while corporate planners struggle with the question of what it is all for. Technical development does not point definitively toward any particular path. Instead, it opens branches, and the final determination of the "right" branch is not within the competence of engineering because it is simply not inscribed in the nature of the technology.

I have studied a particularly clear example of the complexity of the relation between the technical function and meaning of the computer in the case of French videotex. Called "Teletel," this system was designed to bring France into the Information Age by giving telephone subscribers access to databases through a standard dumb terminal. Fearing that consumers would reject anything resembling office equipment, the telephone company attempted to redefine the computer's social image; it was no longer to appear as a filing and calculating device for professionals but was to become a public informational network.

The telephone company designed a new type of terminal, the Minitel, to look and feel like an adjunct to the domestic telephone. The telephonic disguise suggested to some users that they ought to be able to talk to each other on the network. Soon the Minitel underwent a further redefinition at the hands of these users, many of whom employed it primarily for anonymous online chatting with other users in the search for amusement, companionship, and sex.

Thus the design of the Minitel invited communications applications that the company's engineers had not intended when they set about improving the flow of information in French society. Those applications, in turn, connoted the Minitel as a means of personal encounter, the very opposite of the rationalistic project for which it was originally created. The "cold" computer became a "hot" new medium.

At issue in the transformation was not only the computer's narrowly conceived technical function but also the very nature of society it makes possible. Does networking open the doors to the Information Age, where, as rational consumers hungry for data, we pursue strategies of optimization? Or is it a postmodern technology that emerges from the breakdown of institutional and sentimental stability? In this case technology is not merely the servant of some predefined social purpose; it is an environment within which a way of life is elaborated.

In sum, differences in the way social groups interpret and use technical objects are not merely extrinsic but make a difference in the nature of the

objects themselves. *What* the object *is* for the groups that ultimately decide its fate determines what it *becomes* as it is redesigned and improved over time. If this is true, then we can understand technological development only by studying its meaning for the various groups that influence it.

TECHNOLOGICAL HEGEMONY

In addition to the assumptions about individual technical objects we have been discussing so far, meanings belonging to the cultural horizon of society also shape technologies. This second hermeneutic dimension of technology is the basis of modern forms of social hegemony; it is particularly relevant to our original question concerning the inevitability of hierarchy in technological society.

As I will use the term, hegemony is a form of domination so deeply rooted in social life that it seems natural to those it dominates. One might also define it as that aspect of the distribution of social power that has the force of culture behind it.

The term "horizon" refers to culturally general assumptions that form the unquestioned background to every aspect of life. Some of these support the prevailing hegemony. For example, in feudal societies, the "chain of being" established hierarchy in the fabric of God's universe and protected the caste relations of the society from challenge. Under this horizon, peasants revolted in the name of the king, the only imaginable source of power. Rationalization is our modern horizon, and technological design is the key to its effectiveness as the basis of modern hegemonies.

Technological development is constrained by cultural norms originating in economics, ideology, religion and tradition. I discussed earlier how assumptions about the age composition of the labor force entered into the design of nineteenth-century production technology. Such assumptions seem so natural and obvious they often lie below the threshold of conscious awareness.

This is the point of Herbert Marcuse's rationalization (Marcuse 1968). Marcuse shows that Weber confounds the control of labor by management with the control of nature by technology. The search for control of nature is generic, but management arises only against a specific social background, the capitalist system. Workers have no immediate interest in output in this system since their wage is not essentially linked to the income of the firm. Control of human beings becomes all-important in this context. Another way to put it would be to say that top down management is "rational" under the horizon of capitalism, but Weber left off the qualifying phrase.

Through mechanization, some of the control functions are eventually transferred from human overseers and parcelized work practices to machines. Machine design is thus socially relative in a way that Weber never recognized, and the "technological rationality" it embodies is not universal but particular to capitalism. In fact, it is the horizon of all the existing industrial societies, communist as well as capitalist, insofar as they are managed from above.

If Marcuse is right, it ought to be possible to trace the impress of class relations in the very design of production technology, as has indeed been shown by such Marxist students of the labor process as Harry Braverman and David Noble (Braverman 1974; Noble 1984). The assembly line offers a particularly clear instance because its design achieves traditional management goals, such as deskilling and pacing work. Its technologically enforced labor discipline increases productivity and profits by increasing control. However, the assembly line appears as technical progress only in a specific social context. It would not be perceived as an advance in an economy based on workers' cooperatives in which labor discipline was largely self-imposed rather than imposed from above. In such a society, a different technological rationality would dictate different ways of increasing productivity.

This example shows that technological rationality is not merely a belief, an ideology, but is effectively incorporated into the structure of

factors operative in the prevailing rationality. The fact that the argument for the social relativity of modern technology originated in a Marxist context has obscured its most radical implications. We are not dealing here with a mere critique of the property system but have extended the critique down into the technical "base," the forces of production. This approach goes well beyond the old economic distinction between capitalism and socialism, market and plan. Instead, one arrives at a very different distinction between societies in which power rests on the technical mediation of social activities and those that democratize technical control and, correspondingly, technological design.

DOUBLE ASPECT THEORY

The argument to this point might be summarized as a claim that social meaning and functional rationality are inextricably intertwined dimensions of technology. They are not ontologically distinct with meaning in the observer's mind and rationality in the technology proper. Rather they are "double aspects" of the same underlying technical object, each aspect revealed by a specific contextualization.

Functional rationality isolates objects from their original context in order to incorporate them into a theoretical system. The institutions that support this procedure, such as laboratories and research and design centers, themselves form a special context with their own practices and links to various social agencies and powers. The notion of "pure" rationality arises when the work of decontextualization is not itself grasped as a social activity reflecting social interests.

Technologies are selected by these interests from among many possible configurations. Guiding the selection process are social codes established by the cultural and political struggles that define the cultural horizon under which the technology will fall. Once introduced, technology offers a material validation of the social order to which it has been preformed. I call this the "bias" of technology: apparently neu-

of a hegemony. The more technology society employs, the more significant is this support.

As Foucault argued in his theory of "power\knowledge," modern forms of oppression are based not so much on false ideologies as on the specific technical "truths" that found and reproduce the dominant hegemony (Foucault 1977). So long as the contingency of the choice of "truth" remains hidden, the deterministic image of a technically justified social order is projected.

The legitimating effectiveness of technology depends on unconsciousness of the cultural horizon under which it was designed. A recontextualizing critique of technology can uncover that horizon, demystify the illusion of technical necessity, and expose the relativity of the prevailing technical choices. A politics of technology can demand changes reflecting the critique.

The possibility of such a politics is rooted in a peculiar feature of the double aspects of technology. Although function and meaning are analytically distinct aspects of technologies in any temporal cross-section, they interact externally in historical time. They enjoy what might be called a "relationship of double ingression," the data of each invading the other and operating in a paradoxical way on the other's terrain. Everyday experience, the domain of social meaning, is governed by a different logic from the scientific and engineering rationality that presides over the functional logic of technology. Where these contexts are out of alignment, tensions arise that are resolved in the course of history by changes and adjustments in one or both of them.

This is methodologically puzzling but obvious in specific cases. For example, knowledge of risk enters experience as fear or anxiety, that is, an aspect of the meaning of the associated objects. Nuclear power is a case in point. The social meaning of the technology is informed in part by scientific knowledge of risk. But more ancient layers of meaning crystallize around invisible threats and fear of the unknown. Meanwhile, scientists and engineers respond to public perceptions of risk with new designs that promise improved safety.

influences the rational specification of the device. In other fields such as computing, new functionalities are routinely introduced in response to changes in meaning.

THE SOCIAL RELATIVITY OF EFFICIENCY

These issues appear with particular force in the environmental movement. Many environmentalists argue for technical changes that would protect nature and in the process improve human life as well. Such changes would enhance efficiency in broad terms by reducing harmful and costly side effects of technology. However, this program is very difficult to implement in a capitalist society. There is a tendency to deflect criticism from technological processes to products and people, from a priori prevention to a posteriori clean-up. These preferred strategies are costly and reduce efficiency in the short run. This situation has political consequences.

Reducing side effects and restoring the environment are forms of collective consumption, financed by taxes or higher prices. These approaches dominate public awareness. This is why environmentalism is generally perceived as a cost involving trade-offs and not as a rationalization increasing overall well-being. But in a society obsessed by private consumption, that perception is damning. Economists and businesspeople are fond of explaining the price we must pay in inflation and unemployment for worshipping at Nature's shrine instead of Mammon's. Poverty awaits those who will not adjust their social and political expectations to technological imperatives.

This trade-off approach has environmentalists grasping at straws for a strategy. Some hold out the pious hope that people will turn from economic to spiritual values in the face of the mounting problems of industrial society. Others expect enlightened dictators to impose technological reform on an irrational populace. It is difficult to decide which of these solutions is more improbable, but both are incompatible with basic democratic values (Heilbroner 1975).

The trade-off approach confronts us with dilemmas—environmentally sound technology versus prosperity, workers' satisfaction and control versus productivity, and so forth—where what we need are syntheses. Unless the problems of modern industrialism can be solved in ways that both protect nature and win public support, there is little reason to hope that they will ever be solved. But how can technological reform be reconciled with prosperity when it places a variety of new limits on the economy?

The child labor case shows how apparent dilemmas arise on the boundaries of cultural change, specifically, where the social definition of major technologies is in transition. In such situations, social groups excluded from the original design process articulate their unrepresented interests politically. New values the outsiders believe would enhance their welfare appear as mere ideology to insiders who are adequately represented by the existing designs.

This is a difference of perspective, not of nature. Yet the illusion of essential conflict is renewed whenever major social changes affect technology. At first, satisfying the demands of new groups after the fact has visible costs and, if it is done clumsily, will indeed reduce efficiency until better designs are found. But usually better designs are found, and what appeared to be insuperable obstacles to growth dissolve in the face of technological change.

This situation indicates the essential difference between economic exchange and technique. Exchange is all about trade-offs; more of A means less of B. But the aim of technical advance is precisely to avoid such dilemmas with what Simondon calls "concentrations," elegant designs that optimize several variables at once. A single cleverly conceived mechanism may correspond to many different social demands, one structure to many functions. Design is not a zero-sum economic game but an ambivalent cultural process that serves a multiplicity of values and social groups without necessarily sacrificing efficiency.

THE TECHNICAL CODE

That these conflicts over social control of technology are not new can be seen from the interesting case of the "bursting boilers" (Burke

1972). Steamboat boilers were the first technology regulated in the United States. In the early nineteenth-century the steamboat was a major form of transportation, similar to the automobile or airlines today. The United States was a big country without paved roads but with many rivers and canals, hence the reliance on steamboats. But steamboats frequently blew up when the boilers weakened with age or were pushed too hard. After several particularly murderous accidents in 1816, the city of Philadelphia consulted with experts on how to design safer boilers. This was the first time an American governmental institution interested itself in the problem. In 1837, at the request of Congress, the Franklin Institute issued a detailed report and recommendations based on rigorous study of boiler construction. Congress was tempted to impose a safe boiler code on the industry, but boilermakers and steamboat owners resisted, and the government hesitated to interfere with private property.

It took from that first inquiry in 1816 until 1852 for Congress to pass effective laws regulating the construction of boilers. In that time five thousand people died in steamboat accidents. Is this many casualties or few? Consumers evidently were not too alarmed to travel on the rivers in ever increasing numbers. Understandably, the ship owners interpreted this as a vote of confidence and protested the excessive cost of safer designs. Yet politicians also won votes by demanding safety.

The accident rate fell dramatically once thicker walls and safety valves were mandated. Legislation would hardly have been necessary to achieve this outcome had it been technically determined. But in fact boiler design was relative to a social judgment about safety. That judgment could have been made on strictly market grounds, as the shippers wished, or politically, with differing implications for technical design. In either case, those results *constitute* a proper boiler. What a boiler "is" was thus defined through a long process of political struggle culminating finally in uniform codes issued by the American Society of Mechanical Engineers.

This is an example of how technology adapts to social change. What I call the "technical code" of the object mediates the process. That code responds to the cultural horizon of the society at the level of technical design. Quite down-to-earth technical parameters such as the choice and processing of materials are *socially* specified by the code. The illusion of technical necessity arises from the fact that the code is thus literally "cast in iron," at least in the case of boilers.

The conservative antiregulatory approach is based on an illusion. It forgets that the design process always already incorporates standards of safety and environmental compatibility; similarly, all technologies support some basic level of user or worker initiative. A properly made technical object simply *must* meet these standards to be recognized as such. We do not treat conformity as an expensive add-on but regard it as an intrinsic cost. Raising the standards means altering the definition of the object, not paying a price for an alternative good or ideological value, as the trade-off approach holds.

But what of the much discussed cost/benefit ratio of design changes, such as those mandated by environmental or other similar legislation? These calculations have some application to transitional situations, before technological advances responding to new values fundamentally alter the terms of the problem. But all too often, the results depend on economists' very rough estimates of the monetary value of such things as a day of trout fishing or an asthma attack. If made without prejudice, these estimates may well help to prioritize policy alternatives. But one cannot legitimately generalize from such policy applications to a universal theory of the costs of regulation.

Such fetishism of efficiency ignores our ordinary understanding of the concept, which alone is relevant to social decision making. In that everyday sense, efficiency concerns the narrow range of issues that economic actors routinely address. Unproblematic aspects of technology are not included. In theory one can decompose any technical object and account for each of its elements in terms of the goals it meets, whether it be safety, speed, reliability, and so forth, but in practice no

one is interested in opening the "black box" to see what is inside.

For example, once the boiler code is established, such things as the thickness of a wall or the design of a safety valve appear as essential to the object. The cost of these features is not broken out as the specific "price" of safety and compared unfavorably with a more "efficient" but less secure version of the technology. Violating the code in order to lower costs is a crime, not a trade-off. And since all further progress takes place on the basis of the new safety standard, soon no one looks back to the good old days of cheaper, insecure designs.

Design standards are controversial only while they are in flux. Resolved conflicts over technology are quickly forgotten. Their outcomes, a welter of taken-for-granted technical and legal standards, are embodied in a stable code and form the background against which economic actors manipulate unstabilized aspects of technology in the pursuit of efficiency. The code is not varied in real-world economic calculations but treated as a fixed input.

Anticipating the stabilization of a new code, one can often ignore contemporary arguments that will soon be silenced by the emergence of a new horizon of efficiency calculations. This is what happened with boiler design and child labor; presumably, the current debates on environmentalism will have a similar history, and we will someday mock those who object to cleaner air as a "false principle of humanity" that violates technological imperatives.

Noneconomic values intersect the economy in the technical code. The examples we are dealing with illustrate this point clearly. The legal standards that regulate economic activity have a significant impact on every aspect of our lives. In the child labor case, regulation helped to widen educational opportunities with human consequences that are not merely economic in character. In the riverboat case, Americans chose high levels of security, and boiler design came to reflect that choice. Ultimately, this was no trade-off of one good for another but a noneconomic decision about the value of human life and the responsibilities of government.

Technology is thus not merely a means to an end; technical design standards define major portions of the social environment, such as urban and built spaces, workplaces, medical activities and expectations, life patterns, and so on. The economic significance of technical change often pales beside its wider human implications in framing a way of life. In such cases, regulation defines the cultural framework *of* the economy; it is not just an act *in* the economy....

DEMOCRATIC RATIONALIZATION

For generations, faith in progress was supported by two widely held beliefs: that technical necessity dictates the path of development and that the pursuit of efficiency provides a basis for identifying that path. I have argued here that both these beliefs are false and that, furthermore, they are ideologies employed to justify restrictions on opportunities to participate in decision making in industrial society. I conclude that a reform of technological society can support a broader range of values. Democracy is one of these values.

What does it mean to democratize technology? The problem is not primarily one of legal rights but of initiative and participation. Legal forms may eventually routinize claims that are asserted informally at first, but the forms will remain hollow unless they emerge from the experience and needs of individuals resisting a technocratic hegemony.

That resistance takes many forms, from union struggles over health and safety in nuclear power plants to community struggles over toxic waste disposal to political demands for regulation of reproductive technologies. These movements alert us to the need to take technological externalities into account and demand design changes responsive to the enlarged context revealed in that accounting.

Such technological controversies have become an inescapable feature of contemporary

political life, laying out the parameters for official "technology assessment" (Cambrosio and Limoges 1991; Callon et al. 2009). They prefigure the creation of a new public sphere embracing the technical background of social life and a new style of rationalization that internalizes unaccounted costs born by "nature," that is, some-thing or -body exploitable in the pursuit of profit. Here respect for nature is not antagonistic to technology but opens a new path of development.

As these controversies become commonplace, surprising new forms of resistance and new types of demands emerge. The Minitel example is a model of this new situation. In France, the computer was politicized as soon as the government supplied the general public with a highly rationalistic information system. Users "hacked" the network in which they were enrolled and altered its functioning, introducing human communication on a vast scale where only the centralized distribution of data had been planned. The Internet has also given rise to many such innovative public reactions to technology.

Individuals who are incorporated into these new technical networks have learned to resist through the net itself in order to influence the powers that control it. This is not a contest for wealth or administrative power but a struggle to subvert the technical practices, procedures, and designs structuring everyday life.

It is instructive to compare these cases with the movement of AIDS patients for better medical care. Just as a rationalistic conception of the computer tends to occlude its communicative potentialities, so in medicine caring functions have become mere side effects of treatment, which is itself understood in technical terms. Patients become objects of this technique, more or less "compliant" to management by physicians. The incorporation of thousands of incurably ill AIDS patients into this system destabilized it and exposed it to new challenges (Feenberg 1995, chap. 5; Epstein 1996).

The key issue was access to experimental treatment. Clinical research is one way in which a highly technologized medical system can care for those it cannot yet cure. But until quite recently access to medical experiments has been severely restricted by paternalistic concern for patients' welfare. AIDS patients were able to open up access because the networks of contagion in which they were caught were paralleled by social networks that were already mobilized around gay rights at the time the disease was first diagnosed.

Instead of participating in medicine individually as objects of a technical practice, they challenged it collectively and politically. They "hacked" the medical system and turned it to new purposes. Their struggle represents a counter tendency to the technocratic organization of medicine, an attempt at a recovery of its symbolic dimension and caring functions.

As in the case of the Minitel, it is not obvious how to evaluate this challenge in terms of the customary concept of politics. Nor do these subtle struggles against the growth of silence in technological societies appear significant from the standpoint of the reactionary ideologies that contend noisily with capitalist modernism today. Yet the demand for communication that these movements represent is so fundamental that it can serve as a touchstone for the adequacy of political theories of the technological age.

These resistances, like the environmental movement, challenge the horizon of rationality under which technology is currently designed. Rationalization in our society responds to a particular definition of technology as a means to profit and power. A broader understanding of technology suggests a very different notion of rationalization based on responsibility for the human and natural contexts of technical action. I call this "democratic rationalization" because it requires technological advances that can be made only in opposition to the dominant hegemony. It represents an alternative to both the ongoing celebration of technocracy triumphant and the gloomy Heideggerian counterclaim that "Only a God can save us".

Is democratic rationalization in this sense socialist? There is certainly room for discussion

agenda and the old idea of socialism. I believe there is significant continuity. In socialist theory, workers' lives and dignity stood for the larger contexts that modern technology ignores. The destruction of their minds and bodies on the workplace was viewed as a contingent consequence of capitalist technical design. The implication that socialist societies might design a very different technology under a different cultural horizon was perhaps given only lip service, but at least it was formulated as a goal.

We can make a similar argument today over a wider range of contexts in a broader variety of institutional settings with considerably more urgency. I am inclined to call such a position "socialist" and to hope that in time it can replace the image of socialism projected by the failed communist experiment.

More important than this terminological question is the substantive point. Why has democracy not been extended to technically mediated domains of social life despite a century of struggles? Is it because technology is incompatible with democracy or because it has been used to suppress it? The weight of the argument supports the second conclusion. Technology can deliver more than one type of technological civilization. We have not yet exhausted its democratic potential.

REFERENCES

Bloor, David. 1991. *Knowledge and social imagery.* Chicago: University of Chicago Press.

Braverman, Harry. 1974. *Labor and monopoly capital.* New York: Monthly Review.

Burke, John. 1972. Bursting boilers and the federal power. In *Technology and Culture*, ed. M. Kranzberg and W. Davenport. New York: New American Library.

Callon, Michel, Pierre Lascoumbes, and Yannick Barthe. 2009. *Acting in an uncertain world.* Cambridge, MA: MIT Press.

Cambrosio, Alberto, and Camille Limoges. 1991. Controversies as governing processes in technology assessment. *Technology Analysis and Strategic Management* 3 (4): 377–396.

Cunningham, Frank. 1987. *Democratic Theory and Socialism.* Cambridge, UK: Cambridge University Press.

Ellul, Jacques. 1964. *The technological society.* Trans. J. Wilkerson. New York: Vintage.

Epstein, Steven. 1996. *Impure Science.* Berkeley and Los Angeles: University of California Press.

Feenberg, Andrew. 1995. *Alternative modernity: The technical turn in philosophy and social theory.* Berkeley and Los Angeles: University of California Press.

Foucault, Michel. 1977. *Discipline and Punish.* Trans. A. Sheridan. New York: Pantheon.

Heidegger, Martin. 1977. *The question concerning technology.* Trans. W. Lovett. New York: Harper & Row.

Heilbroner, Robert. 1975. *An inquiry into the human prospect.* New York: Norton.

Latour, Bruno. 1987. *Science in action: How to follow scientists and engineers through society.* Cambridge, MA: Harvard University Press.

Latour, Bruno. 1993. *We have never been modern.* Trans. C. Porter. Cambridge, MA: Harvard University Press.

Marcuse, Herbert. Industrialism and capitalism in the work of Max Weber. In *Negations,* 201–226. Trans. J. Shapiro. Boston: Beacon.

Marx, Karl. 1906 reprint. *Capital.* New York: Modern Library.

Miller, Richard W. 1984. *Analyzing Marx: Morality, power and history.* Princeton, NJ: Princeton University Press.

Noble, David. 1984. *Forces of production.* New York: Oxford University Press.

Pinch, Trevor, and Wiebe Bijker. 1989. The social construction of facts and artifacts: Or how the sociology of science and the sociology of technology might benefit each other. In *The social construction of technological systems,* ed. Wiebe Bijker, Thomas Hughes, and Trevor Pinch, 17–50. Cambridge: MIT Press.

Weber, Max. 1958. *The Protestant ethic and the spirit of capitalism.* Trans. T. Parsons. New York: Scribner's.

1.2.3 TECHNOLOGIES OF HUMILITY: CITIZEN PARTICIPATION IN GOVERNING SCIENCE

SHEILA JASANOFF

What role should citizens have in the formation of science and technology policy? Although technological change carries undoubted benefits for society, it also often imposes significant risks on the citizenry. In light of public concerns about risks such as those associated with global climate change, biotechnology, nanotechnology, genetically modified foods, and a variety of other technological advances that are changing our world, how can citizens play a meaningful role in the evaluation of these technologies? The traditional approach to this problem has delegated the responsibility to evaluate new technologies to expert committees composed of scientists, governmental officials, and corporate experts. But several dramatic failures of adequate oversight suggest that this approach is unsatisfactory and needs to be replaced by one in which citizens play a larger deliberative role. Society's values need to be able to influence the development of new technologies even before their widespread introduction, but how can ordinary citizens make informed judgments about uncertainty and risk in such highly technical issues? The challenge, as Sheila Jasanoff says, is not to replace expertise with citizen participation but to create a culture of technology governance that fosters "more meaningful interaction among policy-makers, scientific experts, corporate producers, and the informed public." By doing this, we can better ensure that our ethical values are brought to bear in effective ways in the deliberative processes through which new technologies are introduced.

FOCUS QUESTIONS

1. What was the basis of Vannevar Bush's "social contract for science" that has governed American science policy since the end of World War II? What developments since that time have placed this model in question?
2. What does Jasanoff mean by "technologies of hubris" and "technologies of humility"? What factors suggest that we need to bring greater humility to bear in the selection of science and technology policies?
3. What are some examples of the "participatory turn" in the formation of science and technology policy? What are some problems with the way this approach has been applied in the past?
4. What four questions should we ask about any enterprise that intends to alter human society?
5. What are the four key social technologies that Jasanoff believes should be brought to bear to improve the process of civic deliberation about science and technology?

KEYWORDS

deviant science, extended peer review, institutional learning, technology policy, transparency

In his prescient 1984 book, the sociologist Charles Perrow forecast a series of "normal accidents" in high-risk technologies. The term applied with precision to events that were strung like dark beads through the later years of the twentieth century—most notably, the 1984 chemical plant disaster in Bhopal, India; the 1986 loss of the *Challenger* shuttle and, in the same year, the nuclear plant accident in Chernobyl, USSR; the contamination of blood supplies with the AIDS virus in Europe and North America; the prolonged crisis over BSE ("mad cow disease") in the United Kingdom; and the U.S. space program's embarrassing, although not life-threatening, mishaps with the Hubble telescope's blurry lens, and several lost and extremely expensive Mars explores. To these we may add the discovery of the ozone hole, climate change, and other environmental disasters as further signs of disrepair. Occurring at different times and in vastly different political environments, these events nonetheless served collective notice that human pretensions of control over technological systems need serious reexamination.

American theorists like Perrow chalked up these failings of technology to avoidable error, especially on the part of large organizations (Clarke 1989, Short and Clarke 1992, Vaughan 1996), but some European analysts suggested a more troubling scenario. Passionately set forth by the German sociologist Ulrich Beck (1992), the thesis of "reflexive modernization" argued that risks are endemic in the way that contemporary societies conduct their technologically intensive business. Scientific and technical advances bring unquestioned benefits, but they also generate new uncertainties and failures, so that doubt continually undermines knowledge and unforeseen consequences confound faith in progress. The risks of modernity, Beck suggested, cut across social lines and operate as a great equalizer of classes. Wealth may increase longevity and improve the quality of life, but it offers no certain protection against the ambient harms of technological societies. This observation was tragically borne out when the collapse of the World Trade Center on September 11, 2001 ended the lives of some 3,000 persons, not discriminating among corporate executives, stock market analysts, computer programmers, secretaries, firefighters, policemen, janitors, and restaurant workers. In many other contexts, however, vulnerability remains closely tied to socioeconomic circumstances, inequalities persist in the ability of groups and individuals to defend themselves against risk.

"Risk," on this account, is not a matter of simple probabilities, to be rationally calculated by experts and avoided in accordance with the cold arithmetic of cost-benefit analysis (Graham and Wiener 1995). Rather, it is part of the modern human condition, woven into the very fabric of progress. The problem we urgently face is how to live well with the knowledge that our societies are inevitably "at risk." Critically important normative questions of risk management cannot be addressed by technical experts with conventional tools of prediction. Such questions determine not only whether we will get sick or die, and under what conditions, but also who will be affected and how we should respond to uncertainty and ignorance. Is it sufficient, for instance, to assess technology's consequences, or must we also seek to evaluate its aims? How should we choose when the values of science appear to conflict with other fundamental values? Has our ability to innovate in some areas run unacceptably ahead of our powers of control? Will some of our most revolutionary technologies increase inequality, promote violence, threaten cultures or harm the environment? And are our institutions, national or supranational, up to the task of governing our dizzying technological capabilities? (Never far from the minds of philosophers and authors of fiction, some of these concerns were also famously

articulated in recent times by Bill Joy, co-founder and chief scientist of Sun Microsystems.)

To answer these questions, the task of managing technologies has to go far beyond the model of "speaking truth to power" that once was thought to link knowledge to political action (Price 1965). According to this template, technical input to policy problems must be developed independently of political influences; the "truth" so generated adequately constrains subsequent exercises of political power. The accidents and troubles of the late twentieth century, however, have called into question the validity of this model: both as a descriptively accurate rendition of ways in which experts relate to policy-makers (Jasanoff 1990), and as a normatively acceptable formula for deploying specialized knowledge within democratic political systems. There is growing awareness that even technical policymaking needs to get more political—or, more accurately, to recognize its political foundations more explicitly. Across a widening range of policy choices, technological cultures must learn to supplement the expert's narrow preoccupation with measuring the risks and benefits of innovation with greater attentiveness to the politics of science and technology.

But how can this expansion in the expert's role be reconciled with well-entrenched understandings of the relations between knowledge and power or expertise and public policy? How should these understandings be modified in response to three decades of research on the social dimensions of science? Can we imagine new institutions, processes, and methods for restoring to the playing field of governance some of the normative and political questions that were too long side-lined in assessing the risks and benefits of technology? And are there structured means for cultivating the social capacity for deliberation and reflection on technological change, much as expert analysis of risks has been cultivated for many decades?

There is a growing need, to this end, for what we may call "technologies of humility." These are methods, or better yet institutionalized habits of thought, that try to come to grips with the ragged fringes of human understanding—the unknown, the uncertain, the ambiguous, and the uncontrollable. Acknowledging the limits of prediction and control, technologies of humility confront "head-on" the normative implications of our lack of perfect foresight. They call for different expert capabilities and different forms of engagement between experts, decision-makers, and the public than were considered needful in the governance structures of high modernity. They require not only the formal mechanisms of participation but also an intellectual environment in which citizens are encouraged to bring their knowledge and critical skills to bear on the resolution of common problems.

THE SOCIAL CONTRACT BETWEEN SCIENCE AND THE STATE

In the United States the need for productive working relations between science and the state was famously articulated not by a social theorist or sociologist of knowledge but by the quintessential technical expert: Vannevar Bush, the distinguished Massachusetts Institute of Technology (MIT) engineer and presidential adviser. Bush foresaw the need for major institutional changes following the intense mobilization of science and technology during the Second World War. In 1945 he produced a report, *Science: The Endless Frontier*, which laid the basis for American policy towards science and technology. Science, in Bush's vision, was to enjoy government patronage in peacetime as in war. Control over the scientific enterprise, however, would be wrested from the military and lodged with the civilian scientific community. Basic research, uncontaminated by industrial application or state ambitions, would thrive in the free air of universities. Scientists would establish the substantive aims as well as the intellectual standards for their research. Bush firmly believed that the bountiful results flowing from scientists' endeavors would be translated into beneficial technologies, contributing to the nation's prosperity and progress. Although his design took years to materialize, and even then

was only imperfectly attained, the U.S. National Science Foundation (NSF) eventually emerged as the primary state funder of basic research. (The creation of the National Institutes of Health [NIH] to sponsor biomedical research divided U.S. science policy in a way not contemplated in Bush's original design. In the recent politics of science, NIH budgets have proved consistently easier to justify than appropriations for other branches of science.) The exchange of government funds and autonomy in return for discoveries, technological innovations and trained personnel came to be known as America's "social contract for science."

Signs of wear and tear in the "social contract" appeared in the 1980s. A spate of highly publicized cases of alleged fraud in science challenged the reliability of peer review and, with it, the underlying assumptions concerning the autonomy of science. The idea of science as a unitary practice also broke down as it became clear that research varies from one context to another, not only across disciplines, but—even more important from a policy standpoint—across institutional settings. It was recognized, in particular, that regulatory science, produced to support governmental efforts to manage risk, was fundamentally different from research driven by scientists' curiosity. At the same time, observers of science in society began questioning whether the categories of basic and applied research held meaning in a world where the production and uses of science were densely connected to each other, as well as to larger social and political consequences (Jasanoff, Markle, Petersen, and Pinch 1995).

Rethinking the relations of science with other social institutions generated three major streams of analysis. The first stream takes the "social contract" essentially for granted but points to its failure to work as its proponents had imagined. Many have criticized science, especially university-based science, for deviating from idealized norms of purity and disinterestedness. Despite (or maybe because of) its simplicity, this critique has seriously threatened the credibility of researchers and their claims to autonomy. Others have tried to replace the dichotomous division of *basic* and *applied* science with more differentiated categories, calling attention to the particularities of science done in different settings to meet different objectives. Still others have sought to respecify from the ground up how scientific knowledge is actually produced. This last line of analysis seeks not so much to correct or refine Vannevar Bush's vision of science as to replace it with a more complex account of how knowledge-making fits into the wider functioning of society.

Deviant Science

Scientific fraud and misconduct appeared on the U.S. policy agenda in the 1980s. Political interest reached a climax with the notorious case of alleged misconduct in an MIT laboratory headed by Nobel laureate biologist David Baltimore. He and his colleagues were exonerated after years of inquiry, which included investigations by Congress and the FBI (Kevles 1998). This and other episodes heightened the tendency for policymakers and the public to suspect that all was not in order in the citadels of basic science and greatly increased federal powers for the supervision of research. Some saw the Baltimore affair as a powerful sign that legislators were no longer content with the old social contract's simple *quid pro quo* of money and autonomy in exchange for technological benefits (Guston 2001). Others, like the science journalist Daniel Greenberg (2001), accused scientists of profiting immoderately from their alliance with the state while failing to exercise moral authority or meaningful influence on policy. American science, at any rate, was asked to justify more explicitly the public money spent on it. A token of the new relationship between science and government came with the reform of NSF's peer review criteria in the 1990s. The Foundation now requires reviewers to assess proposals not only on grounds of technical merit, but also with respect to their wider implications for society—thus according greater prominence to science's social utility. In effect, the fraud investigations of the previous decade opened up other taken-for-granted aspects of scientific autonomy and forced scientists to account for their objectives as well as their honesty.

To these perturbations may be added a steady stream of challenges to the supposed

disinterestedness of academic science. In areas ranging from climate change to biotechnology, critics have charged researchers with having sacrificed their objectivity in exchange for grant money or, worse, equity interests in lucrative start-up companies (Boehmer-Christiansen 1994). These allegations have been especially damaging to biotechnology, because that industry benefits significantly from the rapid transfer of skills and knowledge from universities. Since most western governments are committed to promoting such transfers, biotechnology is caught on the horns of a particular dilemma: how to justify its promises of innovation and progress credibly, when the interests of most scientists are aligned with those of industry, government or, occasionally, public interest advocates.

While financially motivated, pro-industry bias has attracted the most criticism, academic investigators have also come under scrutiny for alleged pro-environment and anti-technology biases. In several cases involving biotechnology—in particular, that of the monarch butterfly study conducted by Cornell University scientist John Losey (1999) in the United States, and Stanley Ewen and Arpad Pusztai's (1999) controversial rat-feeding study in the United Kingdom—industry critics questioned the quality of university-based research and implied that political orientations had prompted premature release or over-interpretation of results. In April 2002 a controversy erupted over an article in *Nature* by a University of California scientist, Ignacio Chapela, who concluded that DNA from genetically modified corn had contaminated native species in Mexico. Philip Campbell, the journal's respected editor, did not retract the paper, but stated that "the evidence available is not sufficient to justify the publication of the original paper," and that readers should "judge the science for themselves" (*Washington Times* 2002). As in the Losey and Ewen and Puzstai cases, critics charged that Chapela's science had been marred by non-scientific considerations. Environmentalists, however, viewed all these episodes as pointing to wholesale deficits in knowledge about the long-term and systemic effects of

Context-Specific Science

The second line of attack on the science–society relationship focuses on the basic-applied distinction. One attempt to break out of that dualism was proposed by Donald Stokes (1997), whose quadrant framework, using Louis Pasteur as the prototype, suggested that "basic" science can be done within highly "applied" contexts. Historians and sociologists of science and technology have long observed that foundational work can be done in connection with applied problems, just as applied problem-solving is often required for resolving theoretical issues (for example, in designing new scientific instruments). To date, formulations based on such findings have been slow to take root in policy cultures.

Another example of the contextualizing approach can be found in the work of Silvio Funtowicz and Jerome Ravetz (1992). They proposed to divide the world of policy-relevant science into three nested circles, each with its own system of quality control: (1) "normal science" (borrowing the term from Thomas Kuhn), for ordinary scientific research; (2) "consultancy science," for the application of available knowledge to well-characterized problems; and (3) "post-normal science," for the highly uncertain, highly contested knowledge needed for many health, safety, and environmental decisions. These authors noted that, while, traditional peer review may be effective within "normal" and even "consultancy" science, the quality of "post-normal" science cannot be assured by standard review processes. Instead, they proposed that work of this nature be subjected to *extended peer review*, involving not only scientists but also the stakeholders affected by the use of science. Put differently, they saw accountability, rather than mere quality control, as the desired objective when science becomes "post-normal." (A problem with this analysis lies in the very term "post-normal science." When scientific conclusions are so closely intertwined with social and normative considerations as in Funtowicz and Ravetz's outermost circle, one may just as well call the "product" by another name, such as "socially relevant knowledge" or "socio-

Sheila Jasanoff's 1990 study of expert advisory committees in the United States provided another perspective on this issue by noting that policy-relevant science (also referred to as "regulatory science")—such as science done for purposes of risk assessment—is often subjected to a special kind of "peer review." Regulatory science is reviewed by multidisciplinary committees rather than by individually selected specialists. The role of such bodies is not only to validate the methods by which risks are identified and investigated, but also to confirm the reliability of the agency's interpretation of the evidence. Frequently, regulatory science confronts the need to set standards for objects or concepts whose very existence was not previously an issue for either science or policy: "fine particulate matter" in air pollution control; the "maximum tolerated dose" (MTD) in bioassays; the "maximally exposed person" in relation to air-borne toxics; or the "best available technology" in programs of environmental regulation. In specifying how such terms should be defined or characterized, advisory committees have to address issues that are technical as well as social, scientific as well as normative, regulatory as well as metaphysical. What *kind* of entity, after all, is a "fine" particulate or a "maximally exposed" person, and by what markers can we recognize them? Studies of regulatory science have shown that the power of advisory bodies to definitively address such issues depends on their probity, representativeness, transparency, and accountability to higher authorities—such as courts and the public. In other words, the credibility of regulatory science rests upon factors that have more to do with democratic accountability than with the quality of science as assessed by peer scientists.

New Modes of Knowledge Production

Going beyond the quality and context dependency of science, some have argued the need to take a fresh look at the structural characteristics of contemporary science in order to make it more socially responsive. Michael Gibbons and his co-authors (1994) concluded that the traditional disciplinary science of Vannevar Bush's "endless frontier" has been largely supplanted by a new mode of knowledge production. The salient properties of this new mode, in their view, include the following:

- Knowledge is increasingly produced in contexts of application (i.e., *all* science is to some extent "applied" science);
- Science is increasingly transdisciplinary—that is, it draws on and integrates empirical and theoretical elements from a variety of fields;
- Knowledge is generated in a wider variety of sites than ever before, not just universities and industry, but also in research centers, consultancies, and think-tanks;
- Participants in science have grown more aware of the social implications of their work (i.e., more "reflexive"), just as publics have become more conscious of the ways in which science and technology affect their interests and values.

The growth of this new mode, as Gibbons et al. note, has necessary implications for quality control. Besides old questions about the intellectual merits of their work, scientists are being asked new questions about its marketability, and its capacity to promote social harmony and welfare.

In other work, Helga Nowotny, Peter Scott, and Michael Gibbons (2001) have grappled with the implications of these changes for knowledge production in public domains. Nowotny et al. propose the concept of "socially robust knowledge" as the solution to problems of conflict and uncertainty. Contextualization, in their view, is the key to producing science for public ends. Science that draws strength from its socially detached position is too frail to meet the pressures placed upon it by contemporary societies. Instead, they imagine forms of knowledge that gain robustness from their very embeddedness in society. The problem, of course, is how to institutionalize poly-centric, interactive, and multipartite processes of knowledge-making within institutions that have worked for decades at keeping expert knowledge away from populism and politics. The question confronting the governance of science is how to bring knowledgeable publics into

the front-end of scientific and technological production—a place from which they have historically been excluded.

THE PARTICIPATORY TURN

Changing modes of scientific research and development provide at least a partial explanation for the current interest in improving public access to expert decision-making. In thinking about research today, policy-makers and the public frequently focus on the accountability of science rather than its quality. As the contexts for science have become more pervasive, dynamic and heterogeneous, concerns about the integrity of peer review have transmuted into demands for greater public involvement in assessing the costs and benefits, as well as the risks and uncertainties, of new technologies. Such demands have arisen with particular urgency in the case of biotechnology, but they are by no means limited to this field.

The pressure for accountability manifests itself in many ways, including demands for greater transparency and participation. One notable example came with U.S. federal legislation in 1998, requiring public access, pursuant to the Freedom of Information Act, to all scientific research generated with public funds (Omnibus Consolidated and Emergency Supplemental Appropriations Act of 1999, P.L. 105-277, 1998). The provision was hastily introduced and scarcely debated. Its sponsor, Senator Richard Shelby (R-Alabama), tacked it on as a last-minute amendment to an omnibus appropriations bill. His immediate objective was to force disclosure of data from a controversial study by the Harvard School of Public Health of the health effects of human exposure to fine particulates. This Six Cities Study provided key justification for the U.S. Environmental Protection Agency's stringent ambient standard for airborne particulate matter, issued in 1997. This sweeping enactment showed that Congress was no longer willing to concede unchecked autonomy to the scientific community in the collection and interpretation of data. Publicly funded science, Congress determined, should be available at all times for public review.

Participatory traditions are less thoroughly institutionalized in European policy-making, but in Europe, too, recent changes in the rules and processes governing expert advice display a growing commitment to involving the public in technically-grounded policy decisions. In announcing the creation of a new Directorate General for Consumer Protection, for example, the European Commission observed in 1997 that, "Consumer confidence in the legislative activities of the EU is conditioned by the *quality and transparency* of the scientific advice and its use on the legislative and control process" (emphasis added). The commitment to greater openness is also evident in the strategies of several new United Kingdom expert bodies, such as the Food Standards Agency, created to restore confidence in the wake of the BSE crisis. Similarly, two major public inquiries—the Phillips Inquiry on BSE and the Smith inquiry on the Harold Shipman murder investigation—set high standards for public access to information through the Internet. All across Europe, opposition to genetically modified foods and crops prompted experiments with diverse forms of public involvement, such as citizen juries, consensus conferences, and referenda (Joss and Durant 1995).

Although admirable, formal participatory opportunities cannot by themselves ensure the democratic and deliberative governance of science. There are, to start with, practical problems. People may not be engaged enough or possess enough specialized knowledge and material resources to take advantage of formal procedures. Participation may occur too late to identify alternatives to dominant or default options; some processes, such as consensus conferences, may be too *ad hoc* or issue-specific to exercise sustained influence on policy. Even timely participation does not necessarily improve decision-making. Empirical research has consistently shown that transparency may exacerbate rather than quell controversy, leading parties to deconstruct each other's positions instead of deliberating effectively. Indeed, the Shelby Amendment reflects one U.S. politician's conviction that compulsory disclosure of data will enable challenges to researchers' own

interpretations of their work. It is in this sense an instrument that can be used for fomenting scientific dissent. By contrast, participation constrained by established formal discourses, such as risk assessment, may not admit novel viewpoints, radical critique, or considerations lying outside the taken-for-granted framing of a problem.

TECHNOLOGIES OF HUMILITY

Participation alone, then, does not answer the problem of how to democratize technological societies. Opening the doors to previously closed expert forums is a necessary step—indeed, it should be seen by now as a standard operating procedure of democratic politics. But the formal mechanisms adopted by national governments are not enough to engage the public effectively in the management of global science and technology. What has to change is the *culture* of governance, nationally as well as internationally, and for this we need to address not only the mechanics but also the substance of participatory politics. The issue, in other words, is no longer whether the public should have a say in technical decisions, but how to promote more meaningful interaction among policy-makers, scientific experts, corporate producers, and the informed public.

The analytic ingenuity of modern states has been directed for many decades toward refining what we may call the "technologies of hubris." To reassure their publics, as well as to keep the wheels of science and industry turning, national governments have developed a series of predictive methods (e.g., risk assessment, cost-benefit analysis, climate modeling) that are designed, on the whole, to facilitate management and control, even in areas of high uncertainty (e.g., Porter 1995). These methods achieve their power through claims of objectivity and a disciplined approach to analysis, but they suffer from three significant limitations. First, they show a kind of peripheral blindness toward uncertainty and ambiguity. Predictive methods focus on the known at the expense of the unknown, producing over-confidence in the accuracy and completeness of the pictures they produce. Well-defined, short-term risks command more attention than indeterminate, long-term ones. At the same time, technical proficiency conveys the impression that analysis is not only rigorous, but complete—in short, that it has adequately taken account of all possible risks. Predictive methods tend in this way to downplay what falls outside their field of vision, and to overstate whatever falls within (Irwin and Wynne 1996).

Second, the technologies of predictive analysis tend to preempt political discussion. Expert analytic frameworks create high entry barriers against legitimate outsider positions that cannot express themselves in terms of the dominant discourse (Winner 1986). Claims of objectivity hide the exercise of judgment, so that the normative presuppositions of studies and models are not subjected to general debate. The boundary work that demarcates the space of "objective" policy analysis is carried out by experts, so that the politics of making demarcations remains locked away from public review and criticism (Jasanoff 1990).

Third, predictive technologies are limited in their capacity to internalize challenges that come from outside their framing assumptions. Techniques develop and grow more sophisticated, to be sure, but not necessarily in ways that revisit the values on which they were founded. For example, techniques for assessing chemical toxicity have become ever more refined, but they continue to rest on the demonstrably faulty assumption that people are exposed to one chemical at a time. Synergistic effects, long-term exposures, and multiple exposures are common in normal life but have tended to be ignored as too messy for analysis. Even in the aftermath of catastrophic failures, modernity's predictive models are often adjusted only to take on board lessons that are compatible with their initial assumptions. When a U.S.-designed chemical factory in Bhopal released the deadly gas methyl isocyanate, killing thousands, the international chemical industry made many improvements in its internal accounting and risk communication practices. But no new methods were developed to assess the risks of technology transfer between radically different cultures of industrial production.

At the beginning of the twenty-first century, the unknown, unspecified and indeterminate aspects of scientific and technological development remain largely unaccounted for in policy-making; treated as beyond reckoning, they escape the discipline of analysis as well as politics. What is lacking is not just the knowledge to help fill the gaps, but the processes and methods for eliciting what the public wants and for using what is already known. To bring these dimensions out of the shadows and into the dynamics of democratic debate, they must first be made concrete and tangible. Scattered and private knowledge has to be amalgamated, perhaps even disciplined, into a dependable civic epistemology. The human and social sciences of previous centuries undertook just such a task of translation. They made visible the social problems of modernity—poverty, unemployment, crime, illness, disease, and, lately, technological risk—often as a prelude to rendering them more manageable, using what I have termed the "technologies of hubris." Today, there is a need for "technologies of humility" to complement the predictive approaches: to make apparent the possibility of unforeseen consequences; to make explicit the normative that lurks within the technical; and to acknowledge from the start the need for plural viewpoints and collective learning. How can these aims be achieved?

From the abundant literature on technological disasters and failures, as well as from studies of risk analysis and policy-relevant science, we can abstract four focal points around which to develop the new technologies of humility. They are *framing*, *vulnerability*, *distribution*, and *learning*. Together, they generate the questions we should ask of almost every human enterprise that intends to alter society: what is the purpose; who will be hurt; who benefits; and how can we know? On all these points, we have good reason to believe that wider public engagement would improve our capacity for analysis and reflection. Participation that pays attention to these four points promises to lead to richer deliberation on the substance of decision-making.

Framing

It is an article of faith in the policy literature that the quality of solutions to perceived social problems depends on the adequacy of their original framing (Schon and Rein 1994). If a problem is framed too narrowly, too broadly, or simply wrongly, then the solution will suffer from the same defects. To take a simple example, a chemical testing policy focused on single chemicals cannot produce knowledge about the environmental health consequences of multiple exposures: the framing of the regulatory issue is more restrictive than the actual distribution of chemical-induced risks, and hence is incapable of delivering the optimal management strategies. Similarly, a belief that violence is genetic may discourage the search for controllable social influences on behavior. A focus on the biology of reproduction may delay or impede effective policies for curbing population growth. When facts are uncertain, disagreements about the appropriate frame are virtually unavoidable and often remain intractable for long periods. Yet, few policy cultures have adopted systematic methods for revisiting the initial framing of issues, despite calls to do so (Stern and Fineberg 1996). Frame analysis thus remains a critically important, though neglected, tool of policymaking.

Vulnerability

Risk analysis treats the "at-risk" human being as a passive agent in the path of potentially disastrous events. In an effort to produce policy-relevant assessments, human populations are often classified into groups (e.g., most susceptible, maximally exposed, genetically predisposed, children or women) that are thought to be differently affected by the hazard in question. Based on physical and biological indicators, these classifications tend to overlook the social foundations of vulnerability and to subordinate individual experiences of risk to aggregate numerical calculations (e.g., Irwin and Wynne 1996). Recent efforts to analyze vulnerability have begun to recognize the importance of socio-economic factors, but assessment methods still take populations rather than individuals as the unit of analysis. These approaches not only

disregard differences within groups but reduce individuals to statistical representations. Such characterizations leave out of the calculus of vulnerability such factors as history, place, and social connectedness, all of which may play crucial roles in determining human resilience. Through participation in the analysis of their vulnerability, ordinary citizens might regain their status as active subjects rather than remain undifferentiated objects in yet another expert discourse.

Distribution

Controversies over such innovations as genetically modified foods and stem cell research have propelled ethics committees to the top of the policy-making ladder in several countries. Frequently, however, these bodies are used as "end-of-pipe" legitimation devices, reassuring the public that normative issues have not been omitted from deliberation. The term "ethics," moreover, does not cover the whole range of social and economic realignments that accompany major technological changes, not their distributive consequences, as technology unfolds across global societies and markets. Attempts to engage systematically with distributive issues in policy processes have not been altogether successful. In Europe, consideration of the "fourth hurdle"—the socio-economic impact of biotechnology—was abandoned after a brief debate. In the United States the congressional Office of Technology Assessment, which arguably had the duty to evaluate socio economic impacts, was dissolved in 1995 (Bimber 1996). President Clinton's 1994 injunction to federal agencies to develop strategies for achieving environmental justice produced few dramatic results (Executive Order 12298, 1994). At the same time, episodes like the rebellion against Monsanto's "terminator gene" demonstrate a deficit in the capacity for ethical analysis in large corporations, whose technological products can fundamentally alter people's lives. Sustained interactions between decision-makers, experts and citizens, starting at the upstream end of research and development, could do much to expose the distributive consequences of innovation.

Learning

Theorists of social and institutional learning have tended to assume that what is "to be learned" is never a part of the problem. A correct, or at least a better, response exists, and the only issue is whether actors are prepared to internalize it. In the real world, however, learning is complicated by many factors. The capacity to learn is constrained by limiting features of the frame within which institutions act. Institutions see only what their discourses and practices permit them to see. Experience, moreover, is polysemic, or subject to many interpretations, no less in policy-making than in literary texts. Even when the fact of failure in a given case is unambiguous, its causes may be open to many different readings. Just as historians disagree over what caused the rise or fall of particular political regimes, so policy-makers may find it impossible to attribute their failures to specific causes. The origins of a problem may look one way to those in power, and quite another way to the marginal or the excluded. Rather than seeking monocausal explanations, then, it would be fruitful to design more avenues through which societies can collectively reflect on the ambiguity of their experiences and to assess the strengths and weaknesses of alternative explanations. Learning, in this modest sense, is a suitable objective of civic deliberation.

CONCLUSION

The enormous growth and success of science and technology during the last century has created difficult contradictions for institutions of governance. As technical activities have become more pervasive and complex, so too has the demand grown for more complete and multivalent evaluations of the costs and benefits of technological progress. It is widely recognized that increased participation and interactive knowledge-making would improve accountability and lead to more credible assessments of science and technology. Such approaches would also be consistent with changes in the modes of knowledge production, which have made science more socially embedded and more closely tied to contexts of application. Yet, modern institutions still operate with

conceptual models that seek to separate science from values and emphasize prediction and control at the expense of reflection and social learning. Not surprisingly, the real world continually produces reminders of the incompleteness of our predictive capacities.

To move public discussion of science and technology in new directions, there is a need for "technologies of humility," complementing the predictive "technologies of hubris" on which we have lavished so much of our past attention. These *social technologies* would give combined attention to substance and process, and stress deliberation as well as analysis. Reversing nearly a century of contrary development, these approaches to decision making would seek to integrate the "can do" orientation of science and engineering with the "should do" questions of ethical and political analysis. They would engage the human subject as an active, imaginative agent in making policy, as well as a source of knowledge, insight, and memory.

BIBLIOGRAPHY

Beck, Ulrich. 1992. *Risk Society Towards a New Modernity.* London: Sage Publications.

Bimber, Bruce. 1996. *The Politics of Expertise in Congress: The Rise and Fall of the Office of Technology Assessment.* Albany: State University of New York Press.

Boehmer-Christiansen, Sonja. 1994. "Global Climate Protection Policy: The Limits of Scientific Advice, Parts 1 and 2." *Global Environmental Change* 4(2) 140–159; 4(3)185–200.

Bush, Vannevar. 1945. *Science: The Endless Frontier.* Washington, DC: U.S. Government Printing Office.

Clarke, Lee. 1989. *Acceptable Risk? Making Decisions in a Toxic Environment.* Berkeley: University of California Press.

Ewen, Stanley W. B., and Arpad Pusztai. 1999. "Effect of Diets Containing Genetically Modified Potatoes Expressing Galanthus Nivalis Lectin on Rat Small Intestine." *Lancet* 354: 1353–1354.

Executive Order 12298. 1994. "Federal Actions to Address Environmental Justice in Minority Populations and Low-Income Populations." Washington, DC: February 11.

Funtowicz, Silvio O., and Jerome R. Ravetz. 1992. "Three Types of Risk Assessment and the Emergence of Post Normal Science." In *Social Theories of Risk*, eds. Sheldon Krimsky and D. Golding. New York: Praeger.

Gibbons, Michael, Camille, Limoges, Helga Nowotny, et al. 1994. *The New Production of Knowledge.* London: Sage Publications.

Graham, John D., Jonathan B. Wiener, eds. 1995. *Risk versus Risk Tradeoffs in Protecting Health and the Environment.* Cambridge, MA: Harvard University Press. A preeminent example of the calculative approach.

Greenberg, Daniel S. 2001. *Science, Money, and Politics Political Triumph and Ethical Erosion.* Chicago: University of Chicago Press.

Guston, David H. 2001. *Between Politics and Science Assuring the Integrity and Productivity of Research.* Cambridge: Cambridge University Press.

Irwin, Alan, and Brian Wynne, eds. 1996. *Misunderstanding Science? The Public Reconstruction of Science and Technology.* Cambridge: Cambridge University Press.

Jasanoff, Sheila. 1990. *The Fifth Branch Science Advisers as Policy-makers.* Cambridge, MA: Harvard University Press.

Jasanoff, Sheila, Gerald E. Markle, James C. Petersen, and Trevor Pinch, eds. 1995. *Handbook of Science and Technology Studies.* Thousand Oaks, CA: Sage Publications.

Joss, Simon, and John Durant, eds. 1995. *Public Participation in Science: The Role of Consensus Conferences in Europe.* London: Science Museum.

Kevles, Daniel J. 1998. *The Baltimore Case: A Trial of Politics, Science, and Character.* New York: Norton.

Losey, John E.; L. S. Rayor; and M. E. Carter. 1999. "Transgenic pollen harms monarch larvae." *Nature* 399: 214.

Nowotny, Helga; Peter Scott; and Michael Gibbons. 2001. *Re-Thinking Science: Knowledge and the Public in an Age of Uncertainty.* Cambridge, UK: Polity.

Omnibus Consolidated and Emergency Supplemental Appropriations Act of 1999, P.L. 105–277 (codified at 42 U.S.C. § 1395 et. seq.). 1998. The Office of Management and Budget in the Clinton administration controversially narrowed the scope of the law to apply not to *all* publicly-funded research, but only to research actually relied on in policy-making. The issue was not completely resolved as of May 2005.

Perrow, Charles. 1984. *Normal Accidents: Living with High Risk Technologies*. New York: Basic Books.

Porter, Theodore M. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton: Princeton University Press.

Price, Don K. 1965. *The Scientific Estate*. Cambridge, MA: Harvard University Press. The locus classicus of the view of the right relations between knowledge and power.

Schon, Donald A., and Martin Rein. 1994. *Frame/Reflection: Toward the Resolution of Intractable Policy Controversies*. New York: Basic Books.

Short, James F., and Lee Clarke, eds. 1992. *Organizations, Uncertainties, and Risk*. Boulder: Westview Press.

Stern, Paul C., and Harvey V. Fineberg, eds. 1996. *Understanding Risk: Informing Decisions in a Democratic Society*. Washington, DC: National Academy Press.

Stokes, Donald E. 1997. *Pasteur's Quadrant: Basic Science and Technological Innovation*. Washington, DC: Brookings.

Vaughan, Diane. 1996. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago: University of Chicago Press.

Winner, Langdon. 1986. "On Not Hitting the Tar Baby." In his *The Whale and the Reactor*. Chicago: University of Chicago Press.

INTERNET RESOURCES

European Commission. *1997 Communication of the European Commission on Consumer Health and Safety*. COM (97) 183 fin. Available from http://europa.eu.int/comm/food/fs/sc/index_en.html.

Joy, Bill. 2000. "Why the Future Doesn't Need Us." *Wired*, http://www.wired.com/wired/archive/8.04/joy.html.

Washington Times. "Nature Regrets Publication of Corn Study." April 5, 2002. Available from http://www.washingtontimes.com/national/20020405-9384015.htm.

1.2.4 ADDRESSING TECHNOLOGICAL CHANGE: THE CHALLENGE TO SOCIAL THEORY

JUDY WAJCMAN

It is widely believed that technological innovation is the main driver of social change. In this article, Judy Wajcman, professor of sociology at the London School of Economics discusses how contemporary social theory views social change. In an effort to correct this simplistic one-sided "technological determinist" perspective, Wajcman views the relationship between technology and social change as a symbiotic one—technology and society coevolve with each influencing the other. Similarly for the relationship between science and society, recent studies in the history of science confirm that social values and relations are active shapers of scientific inquiry. The many examples provided by Wajcman in this article illustrate how common technological artifacts have evolved as a result of various social factors often in ways not envisioned by their designers. The "use context" of technological artifacts is what gives them their social meaning.

Wajcman explores the contributions of the social studies of science and technology (STS) movement and discusses how it has broadened sociologists' views of

Source: From Judy Wajcman, "Addressing Technological Change," Current Society, May 2002, Vol. 50(3), 347-363. © 2002 Sage Publications, Inc. Reprinted by permission of Sage Publications, Inc.

the relationship of society and technology from technological determinism to one recognizing that "technological change is shaped by the social circumstances within which it takes place." She provides several examples of how familiar artifacts, such as the microwave oven and the VCR, have been shaped by social values of the users of these devices. But she also argues that social scientists should recognize that technological artifacts are more than just the background context of human action; human sociotechnical practices are forms of power that expand and alter the character of human action itself.

She credits feminist scholars as raising the fundamental question of whether the power of men around the world resulted from their monopolization of technology or because technology is "inscribed with gender power relations." Although there are undoubted gender differences, there is nothing natural or inevitable about the ways in which technologies are defined as "masculine." Citing the increasing number of women in scientific and engineering professions and the writings of feminist authors such as Donna Haraway and Sherry Turkle who embrace technology, Wajcman sees the social shaping of computing and biotechnology as means for women's empowerment.

FOCUS QUESTIONS

1. Discuss the "new technical forces" seen by Wajcman as exerting great influence on society. How might these forces impact your life?
2. What are some of the reasons mentioned by Wajcman for women generally being less involved with the development and adoption of various types of technology?
3. Describe the technologies that some contemporary feminist writers think that women should be more involved with and why they think that is important.
4. What role does Wajcman see virtual reality playing in the future, and how might this technology prove to be beneficial?

KEYWORDS

actor-network theory (ANT), cyborg, gender power relations, interpretive flexibility, post-Fordism, postmodernity, Science, Technology, and Society (STS), technological determinism, technological imperative, technoscience

In their "millennial" reflections on the end of the 20th century and the beginning of the 21st, many sociologists see technology as the impetus for the most fundamental of social trends and transformations. Indeed, understanding the role of technologies in the economy and society is now central to social theory. While there are a variety of social theories that proclaim the radical transformation of society all contain, at their core, claims about technological change and its social impact (Kumar, 1995; Amin, 1994). This is as true of the three paradigmatic theories of the transformation that Western societies are undergoing—the theories of the information society, post-Fordism and postmodernity—as it is of more recent theories of globalization. Much emphasis is placed on major new clusters of scientific and technological innovations, particularly the widespread use of information and communication technologies, and the convergence of ways of life around the globe. The increased automation of production and the intensified use of the computer are said to be revolutionizing the economy and the character of employment. In the "information society" or "knowledge economy," the dominant form of work becomes information and knowledge-based. At the same time leisure, education, family relationships and personal identities are seen as moulded by the pressures exerted and opportunities arising from the new technical forces.

The globalization literature similarly postulates that states and societies across the world are

experiencing historically unprecedented change as they try to adapt to a more interconnected but highly uncertain world (Giddens, 1990, 1996; Castells, 1996; Sassen, 1996; Held et al., 1999). Prominence is given to the intensity, extensity and velocity of global flows, interactions and networks embracing all social domains. For example, Manuel Castells (1996) argues that the revolution in information technology is creating a global economy, the product of an interaction between the rise in informational networks and the process of capitalist restructuring. In the "informational mode of development," labor and capital, the central variables of the industrial society, are replaced by information and knowledge. In the resulting "Network Society," the compression of space and time made possible by the new communication technology alters the speed and scope of decisions. Organizations can decentralize and disperse, with high-level decision making remaining in "world cities" while lower-level operations, linked to the center by communication networks, can take place virtually anywhere. Information is the key ingredient of social organization, and flows of messages and images between networks constitute the basic thread of social structure (1996: 477). For Castells, the information age, organized around "the space of flows and timeless time," marks a whole new epoch in the human experience.

SOCIAL THEORY AND TECHNOSCIENCE

These ideas—or ideas like them (for example, Urray, 2000)—are now commonplace in sociology and I foreground them here to illustrate the centrality of technology to contemporary theories of social, cultural and economic change. There are strong echoes of the earlier "post-industrial society" thesis in these accounts, and its tendency to adopt a technologically determinist stance. Castells explicitly builds on theories of post-industrialism, moving beyond a teleological model and giving the analysis a global reach (Bell, 1973; Touraine, 1971). However, while he explicitly attempts to distance himself from technological determinism, he does not entirely succeed. The idea that technology, specifically information and communication technology, is the most important cause of social change permeates his analysis of Network Society. The result is a rather simplistic view of the role of technology in society. In this, Castells is typical of most scholars of the information society who fail to engage with the burgeoning literature in the social studies of science and technology that has developed over the last two decades.[1] I will return to this in the next section of my essay.

The consequences of this failure are not confined to the theoretical domain. These metanarratives about the information age can themselves help to constitute that which they purport to describe and analyse. Technologically determinist understandings of the economy play an increasingly important role in political discourse. Governments everywhere legitimate much of their policy in terms of a technological imperative. Thus access to Internet resources and educational training for the information age are promoted as the fundamental equity issues. As if in response to Bill Gates's famous claim that the Internet would be a powerful force for eliminating barriers of prejudice and inequality, the UK government, for example, has undertaken to put a computer in every classroom as a guarantee of equality in education. There can be no clearer example of the belief in simple technological fixes to complex social problems. Rhetoric about the "digital divide," that between the information rich and the information poor, serves to camouflage pre-existing patterns of social and class inequality. Furthermore, much government discourse about the inevitability of globalization also invokes new information technology as an autonomous and largely unassailable force. In this process, technocratic discourse, globalization and free market economics coalesce into an extremely powerful ideological force (Johnson, 2000).

The idea that technological rationality translates into political rationality has a long lineage. Frankfurt School writers such as Marcuse (1968, 1969) and Habermas (1971) were wise to it before the dawn of the computer age. Indeed, Habermas's famous essay, dedicated to Marcuse for his 70th birthday, was entitled "Technology

and Science as 'Ideology.' " Habermas argued that the legitimating role of market liberalism was replaced by a new modernist, technocratic ideology in which economic growth and social development appeared to be determined by scientific-technical progress. His discussion of the "scientization of politics," whereby political problems are seen as technical ones and the "expert" directs the politician, has continuing relevance for the nature of democratic decision-making today.

Key to the conceptualization of technical rationality in that prescient 1968 essay was the question of whether science is neutral or intrinsically about domination. The context was the argument, central to Frankfurt School thinking, that Western science and technology form a project of control or domination over both nature and humanity. In *One Dimensional Man* (1968), Marcuse argued that because domination is intrinsic to capitalist science and technology, a new society would require a new science and a new technology. This would be based not on mastery, but on the liberation of nature. By contrast, Habermas departed from the central Frankfurt School position to defend a positivist conception of science and technology as ideologically pure, value-free, and above all neutral. He disputes Marcuse's claim that a science or technology fundamentally different from the one we presently have is a real possibility for advanced societies.

Although the nuances of this highly abstract exchange have been lost over time, it foreshadowed much of the debate about the relationship between technology and society that is still central to the discipline of sociology. Habermas's separation of two analytically distinct spheres of technical rationality and substantive rationality (or "system" and "lifeworld"), where the former may extend beyond its appropriate sphere and "colonize" the lifeworld, is a form of implicit technological determinism. It is the position that is most evident in abstract discussions of social theory. However, within the social studies of science and technology, Marcuse's argument that science and technology are socially constituted has had more resonance.

The writings of the Frankfurt School were to influence many within the political movements of the late 1960s and early 1970s to develop a radical critique of science. In campaigns against an abused, militarized, and polluting science, it was argued that science was directed towards profit and warfare. Initially science itself was seen as neutral or value-free and potentially useful as long as it was in the hands of those working for a just society. Gradually, however, the radical science movement developed a Marxist analysis of the class character of science and its links with capitalist methods of production. A revived political economy of science began to argue that the growth and nature of modern science were related to the needs of capitalist society. Increasingly tied to the state and industry, science had become directed towards domination. The ideology of science as neutral was seen as formulations of this position, associated with the radical science movement, was that "science is social relations" (Young, 1977). The point was that the distinction between science and ideology could not be sustained because the dominant social relations of society at large are constitutive of science.

During this same period a radical shift took place in the history, philosophy and sociology of science, which added weight to the view that science could no longer be understood as simply the discovery of reality. Kuhn's *The Structure of Scientific Revolutions* (1970) marked the beginning of what was to become a major new field of study known as the sociology of scientific knowledge. Its central premise is that scientific knowledge, like all other forms of knowledge, is affected at the most profound level by the society in which it is conducted. Numerous historical and contemporary studies of science, and the manner in which inquiry proceeds, highlight the social aspects of scientific knowledge. One major strand of research is concerned [with the demonstration of] how social interests construct knowledge, while another has focused on the deconstruction of the language of science.[3]

Despite the advances that were made through the critique of science in the 1970s, gender-conscious accounts were rare. It was to await second-wave feminism, and the growing body of feminist scholarship that flourished with it, to

and environmental movements all initially saw science and technology as alien and opposed to women's interests. This was particularly a reaction to the way biology and medical science had cast women as different and inferior, and made a case for biologically determined sex roles. Indeed, the continuing fixation in much scientific inquiry with the study of sex difference illustrates the cogency of the feminist critique of biomedicine. Western science was characterized as a masculine project of reason and objectivity, with women relegated to nature rather than culture. Rejecting scientific knowledge as patriarchal knowledge, there were calls for the development of a new science based on women's values.[3]

SOCIAL STUDIES OF TECHNOLOGY

Building on the legacy of this theoretical and political ferment, the social studies of science and technology (STS) developed in the 1970s. Many of us who got involved then had a simple polemical purpose, to shake the stranglehold that a naïve "technological determinism" had on the dominant understanding of the intertwining of society and technology. We were concerned that this view of technology, as an external, autonomous force exerting an influence on society, narrows the possibilities for democratic engagement with technology. It presents a limited set of options: uncritical embracing of technological change, defensive adaptation to it, or simple rejection of it. Against this, STS has its origins in a belief that the content and direction of technological innovation are amenable to sociological analysis and explanation.

Social scientists have increasingly recognized that technological change is itself shaped by the social circumstances within which it takes place. The "new sociology of technology" set out to demonstrate that technological artefacts are socially shaped, not just in their usage, but especially with respect to their design and technical content. Crucially, it rejects the notion that technology is simply the product of rational technical imperatives, that a particular technology will triumph because it is intrinsically the best. Technical reasons are vitally important. But we need to ask why a technical reason was found to be compelling, when it could have been challenged, and what counts as technical superiority in specific circumstances. STS studies show that the generation and implementation of new technologies involve many choices between technical options. A range of social factors affect which of the technical options are selected. These choices shape technologies and, thereby, their social implications. In this way, technology is a socio technical product, patterned by the conditions of its creation and use.

There is now a vast literature, and a variety of social shaping or constructivist approaches, to the social study of technology.[4] A broad sample of the research can be found in journals such as *Science, Technology, & Human Values*; *Social Studies of Science*; *Science as Culture*; and the books series, *Inside Technology*, with MIT Press.[5] As an introduction to the richness of the field, it may be useful at this point to outline the principal concepts that inform it.

The idea of a technological "system" or "network" has been key. Although technological innovation crucially builds on previous technology, it does so not in the form of separate, isolated devices but as part of a whole, as part of a system. An automatic washing machine, say, can work only if integrated into systems of electricity supply, water supply and drainage. A missile, to take another example, is itself part of an ordered system of component parts—warhead, guidance, control, propulsion—and also part of a wider system of launch equipment and command and control networks. The need for a part to integrate into the whole imposes major constraints on how that part should be designed. A technological system is never merely technical: its real-world functioning has technical, economic, organizational, political, and even cultural elements.

Take something you rarely think twice about—the electric refrigerator. We know from historians of technology that once you could choose between an electric refrigerator and a gas refrigerator, both equally effective (Cowan, 1983). General Electric had the financial resources to invest in the development of the electric model, while the manufacturers of gas refrigerators, although they had a product with real advantages

from the consumer's point of view, lacked the resources to develop and market their machine. Economic power, not technical superiority, gave the electric refrigerator the edge over its competitor. However, the design of this kitchen "white good" was also shaped by the post-Second World War spread of single-family houses with correspondingly small-scale appliances. This built environment in turn sustains the cultural ideal of the separation of the public and private domestic spheres.[6] Gender roles and sexual divisions are part of the socio-technical system or network.

This example illustrates the way technological decisions are the result of "heterogeneous engineering," engineering "social" as well as "technical" phenomena: constructing an environment in which favored projects can be seen as viable (Law, 1987). The usual economic explanation, that assumed that firms simply choose technologies that offer the maximum possible rate of profit, has now been the subject of much criticism. In response, some economists utilize the notions of technological trajectory, path-dependence and lock-in to capture the mechanisms through which the evolution of a technology becomes more and more irreversible. The more technologies are adopted and their problems resolved, the better their performance and the greater their adoption. This clearly generates a powerful path-dependence over time, one that marginalizes competing or new technologies.

STS emphasizes, however, that it is not necessarily technical efficiency but rather the contingencies of socio technical circumstances and the play of institutional interests that favor one technology over another (MacKenzie and Wajcman, 1999: 19–21). Indeed, in situations of technical innovation, costs and profits are inherently uncertain; they cannot be taken as given facts. Economic calculations, such as estimating future costs and profits, are affected by the entire way a society is organized. Even within economics, for example in evolutionary economics, markets are beginning to be understood as embedded in social networks (Granovetter, 1985; Nelson and Winter, 1982).

The general point emerges most sharply when we consider the efficient use of labor, apparently a vital issue in technical change. Noble's (1984) classic study of the development of automatically controlled machine tools in postwar USA shows how production technologies can reflect management's need for control over workers. Noble notes that two options existed: "record playback," whereby the machine merely replicated the manual operations of a skilled machinist; and "numerical control" (NC), in which tool movements were controlled by a mathematical program produced by a technician. He shows how the machine-tool suppliers, technologists, and managers in the aerospace companies deliberately suppressed record playback in favor of NC in order to reduce their reliance on the unionized, craft workers. As it happened, however, management found that it needed to retain skilled machinists to operate the new machines effectively. Thus the intentions underlying the technological design, to shift power from the shop floor to the office, were not fully realized.

Long after artefacts leave the industrial laboratory, the process of technological design is still taking place. Take the example of microwave ovens, a direct descendant of military radar technology, developed for food preparation in US navy submarines. When manufacturers first turned their eyes to the domestic market, they conceived of the microwave as a device to reheat prepared food for use by men, especially single men. As a result, it was marketed as a "brown good" and sold next to hi-fi equipment, televisions, and video recorders—goods for leisure and entertainment. This attempt to frame demand was unsuccessful and subsequently both the product and the consumer were reconstituted as a "white good" for housewives who still want to cook (Cockburn and Ormrod, 1993). The way in which women users appropriated this domestic technology was not foreseen by the male managers and engineers who designed it. The finished form of the microwave, which redefined the gendered character of the user, meant that the microwave literally shifted its place in the department store. It now sits alongside washing machines, fridges and freezers as a humdrum domestic appliance.

These cases highlight the divergent requirements and assumptions of technology developers

and users. The making of the microwave is as much a story about the transformation of a quintessentially human activity, cooking, as it is about a technical invention. Technologies are not fixed at the innovation stage but evolve in their implementation and use. The idea of "interpretative flexibility" captures this malleable character of technologies (Pinch and Bijker, 1987). It emphasizes that there is nothing inevitable about the ways technologies evolve. Rather, technological change is a thoroughly contingent and heterogeneous process. Interpretative flexibility refers to the way in which different groups of people involved with a technology can have different understandings of that technology, including different understandings of its technical characteristics. Thus users can radically alter the meanings and deployment of technologies. Pinch and Bijker's focus is not just on the symbolic meaning of technologies but also includes variation in criteria for judging whether a technology "works." For them, machines work because they have been accepted by relevant social groups. As a result, closure or stabilization occurs such that some artefacts become increasingly the dominant forms of the technology. The fact that a machine "works" needs to be explained rather than taken for granted.

This issue goes right to the heart of decisions about the vast techno-science research and development budgets in, for example, military weapons. Think for a moment about the crucial role that testing plays in attempts to justify the Bush administration's missile defence shield. Yet, testing the accuracy of missiles has never been a straightforward empirical matter. MacKenzie's (1990) study of nuclear ballistic missiles reveals the extent to which definitions of accuracy and reliability are constructed rather than being simply factual. For a start, the conditions for peacetime testing are fundamentally different from those under which missiles would have to operate during a war. MacKenzie's point, however, is both more profound and more general than this. He shows that testing inevitably involves a number of differently constructed background assumptions. As a result, no single test is ever accepted by all the parties involved as the ultimate arbiter. Indeed, it is those most closely involved in the scientific work of testing who have the least certainty about their knowledge of missile accuracy figures. The more one looks inside the "black box" of nuclear weapons technology, or any other technological artefact, "the more one realizes that 'the technical' is no clear-cut and simple world of facts insulated from politics" (1990: 381). Whether or not the "Son of Star Wars" works will necessarily be as much a political as a technical judgement.

STS then sees technology and society as bound together inextricably and the traffic between the two as reciprocal. Indeed, since the widespread adoption of "actor-network theory" (ANT), technology and society are no longer seen as separate spheres, influencing each other.[7] Rather, the metaphor of a "heterogeneous network" conveys the view that technology and society are mutually constitutive: both are made of the same stuff—networks linking human beings and non-human entities. The technological, instead of being a sphere separate from society, is part of what makes large-scale society possible. Their most controversial idea, that we cannot deny *a priori* that non-human actors or "actants" can have agency, has helped us to understand the role of technology in producing social life.

In this way, STS has drawn attention to the neglect of technology or materiality in much social theory. Apart from research concerned with the "impact of technology on society," the main focus of sociology has been on *social structure* and *social relations*. Machines, artefacts, things, have generally been treated as background context, rather than dealt with evenhandedly alongside persons, institutions and events.[8] ANT in particular emphasizes the way material resources, artefacts, and technology make society possible. As Latour (1991b) expresses it: technology is society made durable. STS approaches contribute to an understanding of social change by exploring how technologies and new forms of social life are coproduced. To talk of "social relations" as if they were independent of technology is therefore incorrect. Indeed, what we call the social is bound together as much by the technical as by the social. "Social classes, occupational groups,

organizations, professions—all are held in place by intimately linked social and technical means" (Bijker and Law, 1992: 290). Society itself is built along with objects and artefacts. For this reason, sociology cannot afford to ignore what Latour calls the "missing masses" of "mundane artifacts" (1992).

ANT's conception of the non-human as actant serves as a corrective to a rigid conception of social structure. Its analysis of how sociotechnical networks are stabilized or made durable offers new insights into the perennial sociological concern with social order. ANT views society as a *doing* rather than a *being*. The construction of technologies is also a moving relational process achieved in daily social interactions: entities achieve their form as a consequence of their relations with other entities (Law, 1999: 4). This action-orientation is reminiscent of ethnomethodology and chimes well with the current fascination in sociology with "performativity." While ANT is similarly ill equipped to deal with power as a structural phenomenon, it does sensitize us to the materiality of social relations.

One way of expressing this is in terms of the power of objects. The conceptions of power that prevail in the social sciences tend to neglect this form of power (with the important exception of military technology). Writers such as Lukes (1974) and Held (1995) define power in terms of the capacity of social agents, agencies and social institutions. For example, technology does not figure as one of Held's seven key "sites of power," which he defines as an interaction context or institutional milieu through which power operates to mould people's life-chances and effective participation in public decision-making (1995: 173).[9] Even Foucauldian-influenced writers, leaving aside Foucault's classic metaphor of the "panopticon," limit their discussion of power relations to social technologies.

This neglect is not surprising given that when technical systems are completely integrated into the social fabric, they become "naturalized," disappearing into the landscape. Take for example the way seemingly innocuous technologies such as photography and film assume, privilege and construct whiteness. Dyer (1997) describes how it is extremely difficult to film black and white faces in the same film and do equal justice to both. Each requires a completely different handling of lighting, make-up, and film developing. The variation in filming conditions means that when black and white actors are portrayed together, one group tends to lose out, and systematically it is black actors who are technologically short-changed. Dyer traces this bias in the use of film techniques to the film industry's origins in the USA and Europe. From the mid-19th century, experiments with the chemistry of photographic stock, aperture size, length of development and artificial light all proceeded on the assumption that what had to be got right was the look of the white face. By the time of film (some 60 years after the first photographs), technologies and practices were already well established and shaped subsequent uses.

This idea of the power exercised by objects is generalized in Latour's (1992) concept of "delegation to non humans." His popular examples of automatic doors and road bumps show how technical objects define actors, the space in which they move and the ways in which they behave and interact. Fittingly called "sleeping policemen," road bumps are delegated the job of reducing motorists' speed where the rule of law does not suffice. In this way, it can be said that the material world itself exercises a kind of agency.

GENDER AND TECHNOLOGY

Much of the most exciting work in STS is being contributed by feminist scholars, reflecting the growing influence of feminist theory on mainstream sociology. Feminists have identified men's monopoly of technology as an important source of their power; women's traditional lack of technological skills is an important element in their dependence on men. A key issue has been whether the problem lies in men's domination of technology, or whether technology itself is inscribed with gender power relations. A long-standing theme has been the impact or effect of technological change on the sexual division of labor in both paid and unpaid work (see, for example, Cowan, 1983; Hartmann et al., 1986,

1987). Consistent with the STS tradition, feminist technology studies conceive of a two-way mutually shaping relationship between gender and technology in which technology is both a source and consequence of gender relations and vice versa (Wajcman, 1991). More recently, this mutual shaping framework has been "recast in a post-structural trope in which gender and technology are seen as *coproduced*" (Faulkner, 2001: 82; see also Berg, 1996). Such feminist perspectives on technology add a new dimension to sociological analyses of gender difference and sexual inequality.

There is nothing natural and inevitable about the ways in which technology is defined as masculine, and masculinity is identified with technical competence. Indeed, it was only with the formation of engineering as a white, male middle-class profession that "male machines rather than female fabrics" became the modern markers of technology (Oldenziel, 1999). During the 19th century mechanical and civil engineering increasingly came to define what technology is, diminishing the significance of both artefacts and forms of knowledge associated with women. At the same time, femininity was being reinterpreted as incompatible with technological pursuits. The continuing under-representation of women in engineering, and indeed, all scientific and technical institutions, is a legacy of this history.

This conventional understanding of technology, with its masculine connotations, has characterized much of the research literature. Within mainstream STS, the ways in which technological objects may be shaped by the operation of gender interests and gender identities have often been overlooked. While innovations are seen as socio technical networks, it has largely been incumbent on feminists to demonstrate that social relations include gender relations. Despite a burgeoning literature on men and masculinities, the critical role played by technology in hegemonic masculinity has been largely ignored (see, for example, Connell, 1987). This problem was exacerbated in constructivist studies of technology by their concentration on the social groups or networks that directly influence the innovation process. As was commonly assumed that gender interests were not being mobilized. This had the effect of excluding from the picture both the women working further downstream from the design process and those whose primary relation to technology is as users (Star, 1991; Wajcman, 2000).

Indeed, technological systems implicitly place men's experiences and men's investments at the center, without acknowledging their specificity. The corollary is the simultaneous denial of other realities such as women's. Feminists have demonstrated that the marginalization of women from the technological community has a profound influence on the design, technical content, and use of artefacts (Berg, 1996; Cockburn, 1983, 1985; Grint and Gill, 1995). This is not meant to imply an essentialist or dualistic analysis in which men and women are treated as homogeneous groups. Rather, a gender analysis shows how preferences for different technologies are shaped by a set of social arrangements that reflect men's power and resources in the wider society.

To what extent, for example, are the needs and interests of women as users taken on board in the process of designing ICTs for the home of the future? Anne-Jorunn Berg's (1994) study of the prototypes for "smart" houses reveals a preoccupation with innovations that automate control of leisure technologies and heating, lighting and security. The core aim of automating the smart house turns out to be "integration," that is, centralized control and regulation of all functions in a local network or "house-brain." The designers and producers of the technological home show little interest in housework. It is the technology as such, the way artefacts function in technical terms, that fascinates the designers. The target consumer is the technically interested man, someone in his own image. The neglect of women's knowledge, experience and skills as a resource for technical innovation in the home is evidence of a profoundly gendered innovation process. Berg argues that many innovations that would simplify some tasks, like cleaning and ironing, have not been developed. While there may be limits to the mechanization of household labor, it seems that a labor-saving "self-cleaning" house

Gender relations then are an integral constituent of the institutions and projects from which technologies emerge. As a result, technological artefacts can be encoded by gender, both materially and symbolically. This is not to imply, however, that gender is a fixed and unitary phenomenon, which exists prior to and independently of technology, and then becomes embedded within it. Contemporary feminist theory stresses that gender is not fixed in advance of social interaction, but is constructed in interaction. Individuals act or perform gender, and demonstrate their gender identity. Gender is a social achievement. This notion of "gender as doing" fits well with the action-orientation of constructivist studies of technology. Both technology and gender are now conceptualized as fluid, relational processes. Thus masculinities and femininities are constituted simultaneously with the production and consumption of technologies. The key argument here is that our relationship to technology is pivotal to the discourse of gender dualisms and gender difference. This issue is explored in the abundant scholarship on how technology as culture is implicated in the construction of subject identity for both sexes.

The most influential feminist commentator writing in this vein is Donna Haraway (1997). She argues that we should embrace the positive potential of technoscience and is sharply critical of those who reject technology. Famously, she prefers to be a "cyborg"—a hybrid of organism and machine parts—rather than an ecofeminist "goddess." She notes the great power of science and technology to create new meanings and new entities, to make new worlds. Genetic engineering, reproductive technology, and the advent of virtual reality are all seen as fundamentally affecting the basic categories of "self" and "gender." She positively revels in the very difficulty of predicting what technology's effects will be and warns against any purist rejection of the "unnatural," hybrid, entities produced by biotechnology. The cyborg metaphor has been widely adopted, bridging "the language of material feminists working on issues of gender and technoscience, and postmodern feminists working with cultural studies and textual deconstruction" (Kirkup et al., 2000: 4–5). Haraway's work encourages a critical optimism in feminist analyses and politics, exploring the ways in which women's everyday lives are entwined with technologies, albeit in diverse and often complex and contradictory ways.

Although they come from somewhat different intellectual traditions, there is a synchronicity between Haraway's depiction of our cyborg-like existence and STS's emphasis on the seamless web of the socio-technical. This is particularly evident in the increasing preoccupation in sociological theory with the body, sexuality and the role of biomedical technologies—technologies for the body. In studies of childbirth and contraception, in vitro fertilization, cosmetic surgery and genetic engineering, feminists argue that there is no such thing as the natural, physiological body. One consequence of this work is that the conventional distinction between sex (natural) and gender (social) has been thoroughly contested and deconstructed. Technologies, like science, are now seen as contributing to the stabilization of meanings of the body. With the rise of modern science, bodies have become objects that can be transformed with an increasing number of tools and techniques. Modern bodies are made and remade through science and technology; they too are technological artefacts. Life itself (human, plant, and animal) has been biomedicalized (Clarke, 1998).

These studies illustrate the connection between the gendered discourses of the biomedical sciences and the institutionalization of medical techniques applied to women's bodies. They remind us that the conceptualization of male and female bodies as essentially different, rather than similar, is a modern one, dating only from the 18th century. The identification of the female body as the Other resulted in positioning it as the quintessential medical object. Women were defined in terms of sex and reproduction, and this was reflected in the establishment of gynaecology as a separate branch of medicine. With the emergence of sex endocrinology in the 1920s and 1930s, the essence of femininity came to be located in hormones (Oudshoorn, 1994). The development of a contraceptive pill for women, rather than for men, was the product not of

biological necessity but of 20th-century biomedical practices.

In the same way as biomedical technologies are seen as transforming the relations between the self, the body and machines, so too are information and communication technologies. There is a profusion of studies on the Internet, cyberspace and virtual reality. Many authors celebrate the myriad ways modern technology is challenging traditional notions of gender identity and creating new meanings and new entities. In *Life on the Screen: Identity in the Age of the Internet*, for example, Sherry Turkle (1995: 12) enthuses about the potential for people "to express multiple and often unexplored aspects of the self, to play with their identity and to try out new ones." It is the increasingly interactive and creative nature of computing technology that now enables millions of people to live a significant segment of their lives in virtual reality. Moreover, it is in this computer-mediated world that people experience a new sense of self that is decentered, multiple and fluid. In this respect, she argues, the Internet is the material expression of the philosophy of postmodernism.

New communication technologies have certainly opened up novel possibilities for sociality in cyberspace. Excitement has been generated about the capacity they provide to create new communities and the new tools they offer for global information exchange and political organizing, as well as all forms of commerce. Indeed, that communities can exist entirely in cyberspace, without face-to-face interaction, challenges classical sociological theories of "community." Virtual communities reflect the human need for communication and belonging without imposing any of the traditional constraints, duties and responsibilities. Of course the forms they take are not as transcendent as we might imagine—so for example the novel forms of technoculture articulated in cyberspace are often characterized by violence and pornography. This is a reminder that while the pioneering possibilities of digitalization, the Internet and cyberspace may be thrilling, digital environments largely simulate the mundane outside world. Revolutions in technology do not create new societies, but they do change the terms in which social, political and economic relations are played out.

CONCLUSION

I began this article with a brief discussion of some recent contributions to social theory that claim technological change is shaping new social forms in the late modern social world. For many of these writers, these developments represent significant disruptions of past expectations and herald an increasingly complex and fragmented world that requires new social theories. Yet, I have suggested that these theoretical arguments are much less sophisticated than they initially appear and frequently resort to an implicit technological determinism. In contrast, I have argued that the field of social studies of science and technology provides rich resources for understanding technology as a socio-technical product, patterned by the relations of its production and use.

For many commentators, the present condition of social theory is one of increasingly discordant claims and fragmentation. Yet much of this social theory is disconnected from the social research that would give much greater insight into the nature of technical change and its social character. We do now have a considerable legacy of sophisticated theoretical writing and detailed case studies around which it has been produced. It is work that deserves greater purchase in the public debate. From my review of the field of STS, sociology looks like a healthy collective endeavor, with much evidence of productive convergence where arguments are shaped by their mutual engagement.

NOTES

1. This technological determinism is also evident in Giddens; see, for example, *The Consequences of Modernity* (1990).
2. This project has been particularly associated with the Edinburgh School, and the "strong" programme in the sociology of knowledge. For an introduction to this literature, see Barnes and Edge (1982) and Knorr-Cetina and Mulkay (1983).
3. For an excellent map of the field of gender and science, see Harding (1986).

4. Whereas references to the "new sociology of technology" were common in the 1980s, the term "constructivist studies of technology" is often now used to include the actor-network approach, the social-constructivist approach, the social shaping approach and the systems approach to technology studies. See Bijker (1995), Bijker et al. (1987), Bijker and Law (1992), Callon (1986), Latour (1987, 1991a, 1993), Law and Hassard (1999), and MacKenzie and Wajcman (1999). For a recent overview of the field, see Part One of the revised edition of MacKenzie and Wajcman (1999). See also the Handbook produced by the Society for the Social Studies of Science, Jasanoff et al. (1994).
5. The journal *Technology and Culture* contains some of the best case studies in the social history of technology.
6. See Chapter 5 of Wajcman (1991) for an elaboration of this point.
7. See Callon (1986), Latour (1987, 1991b, 1993), and Law and Hassard (1999).
8. Interestingly, during the same period, the concept of "material culture" has been restored to the core concerns of social anthropology, as illustrated by the 1996 launch of *Journal of Material Culture* (London: Sage).
9. The seven sites are the body, welfare, culture, civic associations, the economy, violence and coercive relations, regulatory and legal institutions.

REFERENCES

Amin, A., ed. (1994) *Post-Fordism: A Reader*. Oxford: Blackwell.

Barnes, B. and Edge, D., eds. (1982) *Science in Context Readings in the Sociology of Science*. Buckingham: Open University Press.

Bell, D. (1973) *The Coming of Post-Industrial Society*. New York: Basic Books.

Berg, A. J. (1994) "A Gendered Socio-technical Construction: The Smart House," in C. Cockburn and R. F. Dilic (eds.) *Bringing Technology Home: Gender and Technology in a Changing Europe*. Milton Keynes: Open University Press.

Berg, A. J. (1996) *Digital Feminism*, Report No. 28. Dragvoll, Norway: Senter for Teknologi og Samfunn, Norwegian University of Science and Technology.

Bijker, W. E. (1995) *Of Bicycles, Bakelites, and Bulbs Toward a Theory of Sociotechnical Change*. Cambridge: MA: MIT Press.

Bijker, W. E. and Law, J., eds. (1992) *Shaping Technology/Building Society Studies in Sociotechnical*

Bijker, W. E., Hughes, T. P., and Pinch, T., eds. (1987) *The Social Construction of Technological Systems*. Cambridge, MA: MIT Press.

Callon, M. (1986) "Some Elements of a Sociology of Translation Domestication of the Scallops and the Fisherman of St Brieue Bay," in J. Law (ed.) *Power, Action and Belief: A New Sociology of Knowledge?*, pp. 196–233. London: Routledge.

Castells, M. (1996) *The Rise of the Network Society*. Oxford: Blackwell.

Clarke, A. (1998) *Disciplining Reproduction: Modernity, American Life Sciences, and "the Problems of Sex."* Berkeley: University of California Press.

Cockburn, C. (1983) *Brothers: Male Dominance and Technological Change*. London: Pluto Press.

Cockburn, C. (1985) *Machinery of Dominance: Women, Men and Technical Know-How*. London: Pluto Press.

Cockburn, C. and Ormrod, S. (1993) *Gender and Technology in the Making*. London: Sage.

Connell, B. (1987) *Gender and Power*. Cambridge: Polity Press.

Cowan, R. S. (1983) *More Work for Mother: The Ironies of Household Technology from the Open Hearth to the Microwave*. New York: Basic Books.

Dyer, R. (1997) *White*. London: Routledge.

Faulkner, W. (2001) "The Technology Question in Feminism: A View from Feminist Technology Studies," *Women's Studies International Forum* 24(1): 79–95.

Giddens, A. (1990) *The Consequences of Modernity*. Cambridge: Polity Press.

Giddens, A. (1996) "Globalization: A Keynote Address," *UNRISD News* 15.

Granovetter, M. (1985) "Economic Action and Social Structure: The Problem of Embeddedness," *American Journal of Sociology* 91(3): 481–510.

Grint, K. and Gill., R., eds. (1995) *The Gender-Technology Relation: Contemporary Theory and Research*. London: Taylor & Francis.

Habermas, J. (1971) "Technology and Science as 'Ideology,' " in *Toward a Rational Society*. London: Heinemann.

Haraway, D. (1997) *Modest Witness@Second Millennium*. New York: Routledge.

Harding, S. (1986) *The Science Question in Feminism*. New York: Cornell University Press.

Hartmann, H., Kraut, R., and Tilly, L., eds. (1986, 1987) *Computer Chips and Paper Clips: Technology and Women's Employment*, Volumes 1 and 2. Washington, DC: National Academy Press.

Held, D. (1995) *Democracy and the Global Order*.

Held, D., McGrew, A., Goldblatt, D., and Perraton, J. (1999) *Global Transformation: Politics, Economics and Culture*. Cambridge: Polity Press.

Jasanoff, S., Markle, G., Petersen, J., and Pinch, T., eds. (1994) *Handbook of Science and Technology Studies*. Thousand Oaks, CA: Sage.

Johnson, C. (2000) *Governing Change: From Keating to Howard*. Brisbane: University of Queensland Press.

Kirkup, G., Janes, L., Woodward, K., and Hovenden, F., eds. (2000) *The Gendered Cyborg: A Reader*. London: Routledge.

Knorr-Cetina, K. and Mulkay, M., eds. (1983) *Science Observed Perspectives in the Social Studies of Science*. London: Sage.

Kuhn, T. S. (1970) *The Structure of Scientific Revolutions*. Chicago, IL: The University of Chicago Press.

Kumar, K. (1995) *From Post-Industrial to Post-Modern Society: New Theories of the Contemporary World*. Oxford: Blackwell.

Latour, B. (1987) *Science in Action*. Buckingham: Open University Press.

Latour, B. (1991a) *The Pasteurization of France*. Cambridge, MA: Harvard University Press.

Latour, B. (1991b) "Technology is Society Made Durable," in J. Law (ed.) *A Sociology of Monsters Essays on Power, Technology and Domination*, pp. 103–31. London: Routledge.

Latour, B. (1992) "Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts," in W. Bijker and J. Law (eds.) *Shaping Technology/Building Society Studies in Sociotechnical Change*. Cambridge, MA: MIT Press.

Latour, B. (1993) *We Have Never Been Modern*. Hemel: Hempstead Harvester.

Law, J. (1987) "Technology and Heterogeneous Engineering: The Case of Portuguese Expansion," in W. E. Bijker, T. P. Hughes, and T. Pinch (eds.) *The Social Construction of Technological Systems*. Cambridge, MA: MIT Press.

Law, J. (1999) "After ANT Complexity, Naming and Topology," in J. Law and J. Hassard (eds.) *Actor Network Theory and After*, pp. 1–14. Oxford: Blackwell.

Law, J. and Hassard, J., eds. (1999) *Actor Network Theory and After*. Oxford: Blackwell.

Lukes, S. (1974) *Power*. London: Macmillan.

MacKenzie, D. (1990) *Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance*. Cambridge, MA: MIT Press.

MacKenzie, D. (1996a) *Knowing Machines: Essays on Technical Change*. Cambridge, MA: MIT Press.

MacKenzie, D. (1996b) "How Do We Know the Properties of Artefacts' Applying the Sociology of Knowledge to Technology," in R. Fox (ed.) *Technological Change: Methods and Themes in the History of Technology*. London: Harwood.

MacKenzie, D. and Wajcman, J., eds. (1999) *The Social Shaping of Technology*, rev and expanded edn. Buckingham: Open University Press. (Orig. pub. 1985.)

Marcuse, H. (1968) *One Dimensional Man*. London: Sphere.

Marcuse, H. (1969) *An Essay on Liberation*. Boston, MA: Beacon.

Nelson, R. R. and Winter, S. G. (1982) *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press.

Noble, D. F. (1984) *Forces of Production. A Social History of Industrial Automation*. New York: Knopf.

Oldenziel, R. (1999) *Making Technology Masculine: Men, Women and Modern Machines in America*. Amsterdam: Amsterdam University Press.

Oudshoorn, N. (1994) *Beyond the Natural Body: An Archeology of Sex Hormones*. London: Routledge.

Pinch, T. and Bijker, W. (1987) "The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other," in W. E. Bijker, T. P. Hughes and T. Pinch (eds.) *The Social Construction of Technological Systems*. Cambridge, MA: MIT Press.

Sassen, S. (1996) *Losing Control? Sovereignty in an Age of Globalization*. New York: Columbia University Press.

Star, S. L. (1991) "Power, Technologies and the Phenomenology of Conventions: On Being Allergic to Onions," in J. Law (ed.) *A Sociology of Monsters Essays on Power, Technology and Domination*, pp. 26–56. London: Routledge.

Touraine, A. (1971) *The Post-Industrial Society*. London: Routledge.

Turkle, S. (1995) *Life on the Screen: Identity in the Age of the Internet*. New York: Simon & Schuster.

Urray, J. (2000) "Mobile Sociology," *British Journal of Sociology* 51(1): 185–203.

Wajcman, J. (1991) *Feminism Confronts Technology*. Pennsylvania: Penn State University Press and Cambridge: Polity Press.

Wajcman, J. (2000) "Reflections on Gender and Technology Studies: In What State Is the Art?," *Social Studies of Science* 30(3): 447–64.

Young, B. (1977) "Science *Is* Social Relations," *Radical Science Journal* (5): 65–129.

1.3 ETHICAL PERSPECTIVES

1.3.1 TECHNOLOGY AND RESPONSIBILITY: REFLECTIONS ON THE NEW TASK OF ETHICS

HANS JONAS

Hans Jonas is one of a relatively small number of twentieth-century philosophers who have reflected carefully on the relationship between technology and ethics. He is also known for his writing on topics in biomedical ethics. He ended his professional career as Alvin Johnson Professor Emeritus of Philosophy at the New School for Social Research.

In this essay that originally appeared in *Social Research* in 1973, Jonas develops the view that traditional ethics is incapable of handling the kinds of ethical problems that are being created by our contemporary global technological civilization. His view is that to adequately evaluate the ethical significance of contemporary science and technology, we need to do so through the lens of a new kind of ethics, what he terms an "ethics of responsibility." He subsequently expanded the ideas presented in this essay in a book published in 1984 entitled *The Imperative of Responsibility: In Search of an Ethics for the Technological Age*.

FOCUS QUESTIONS

1. What are the three characteristics of traditional "neighbor ethics" that Jonas isolates, and in what respects has modern technology made traditional ethics obsolete?
2. How does Jonas's view that "man himself has been added to the objects of technology" compare to the discussion of the ethics of human cloning by Leon Kass (Selection 2.5.3) and Michael Sandel (Selection 2.5.4)?
3. What do you think Jonas means when he says, "We need wisdom the most when we believe in it the least"?

KEYWORDS

behavior control, environmental ethics, genetic engineering, morality, Utopia

All previous ethics—whether in the form of issuing direct enjoinders to do and not to do certain things, or in the form of defining principles for such enjoinders, or in the form of establishing the ground of obligation for obeying such principles—had these interconnected tacit premises in common: that the human condition, determined by the nature of man and the nature of things, was given once for all; that the human good on that basis was readily determinable; and that the range

Source: From *Philosophical Essays* by Hans Jonas. Copyright © 1974 by Prentice-Hall. Reprinted by permission of Eleanore Jonas.

of human action and therefore responsibility was narrowly circumscribed. It will be the burden of my argument to show that these premises no longer hold, and to reflect on the meaning of this fact for our moral condition. More specifically, it will be my contention that with certain developments of our powers the *nature of human action* has changed, and since ethics is concerned with action, it should follow that the changed nature of human action calls for a change in ethics as well: this not merely in the sense that new objects of action have added to the case material on which received rules of conduct are to be applied, but in the more radical sense that the qualitatively novel nature of certain of our actions has opened up a whole new dimension of ethical relevance for which there is no precedent in the standards and canons of traditional ethics.

I

The novel powers I have in mind are, of course, those of modern *technology*. My first point, accordingly, is to ask how this technology affects the nature of our acting, in what ways it makes acting under its dominion *different* from what it has been through the ages. Since throughout those ages man was never without technology, the question involves the human difference of *modern* from previous technology. Let us start with an ancient voice on man's powers and deeds which in an archetypal sense itself strikes, as it were, a technological note—the famous Chorus from Sophocles' *Antigone*.

> *Many the wonders but nothing more wondrous than man.*
> *This thing crosses the sea in the winter's storm, making his path through the roaring waves.*
> *And she, the greatest of gods, the Earth—deathless she is, and unwearied—he wears her away as the ploughs go up and down from year to year and his mules turn up the soil.*
> *The tribes of the lightbearted birds he ensnares, and the races of all the wild beasts and the salty brood of the sea, with the twisted mesh of his nets, he leads captive, this clever man.*
> *He controls with craft the beasts of the open air, who roam the hills. The horse with his shaggy mane he holds and harnesses, yoked about the neck, and the strong bull of the mountain.*
> *Speech and thought like the wind and the feelings that make the town, he has taught himself, and shelter against the cold, refuge from rain. Ever resourceful is he.*
> *Tie faces no future helpless. Only against death shall he call for aid in vain. But from baffling maladies has he contrived escape.*
> *Clever beyond all dreams the inventive craft that he has which may drive him one time or another to well or ill.*
> *When he honors the laws of the land the gods' sworn right high indeed in his city; but stateless the man who dares to do what is shameful.*

This awestruck homage to man's powers tells of his violent and violating irruption into the cosmic order, the self-assertive invasion of nature's various domains by his restless cleverness; but also of his building—through the self-taught powers of speech and thought and social sentiment—the home for his very humanity, the artifact of the city. The raping of nature and the civilizing of himself go hand in hand. Both are in defiance of the elements, the one by venturing into them and overpowering their creatures, the other by securing an enclave against them in the shelter of the city and its laws. Man is the maker of his life *qua* human, bending circumstances to his will and needs, and except against death he is never helpless.

Yet there is a subdued and even anxious quality about this appraisal of the marvel that is man, and nobody can mistake it for immodest bragging. With all his boundless resourcefulness, man is still small by the measure of the elements: precisely this makes his sallies into them so daring and allows those elements to tolerate his forwardness. Making free with the denizens of land and sea and air, he yet leaves the encompassing nature of those elements unchanged, and their generative powers undiminished. Them he cannot harm by carving out his little dominion from theirs. They last, while his schemes have their short lived way. Much as he harries Earth, the greatest of gods, year after year with his plow—she is ageless and unwearied; her enduring patience he must

and can trust, and he must conform. And just as ageless is the sea. With all his netting of the salty brood, the spawning ocean is inexhaustible. Nor is it hurt by the plying of ships, nor sullied by what is jettisoned into its deeps. And no matter how many illnesses he contrives to cure, mortality does not bow to cunning.

All this holds because man's inroads into nature, as seen by himself, were essentially superficial, and powerless to upset its appointed balance. Nor is there a hint, in the *Antigone* chorus or anywhere else, that this is only a beginning and that greater things of artifice and power are yet to come—that man is embarked on an endless course of conquest. He had gone thus far in reducing necessity, had learned by his wits to wrest that much from it for the humanity of his life, and there he could stop. The room he had thus made was filled by the city of men—meant to enclose and not to expand—and thereby a new balance was struck within the larger balance of the whole. All the well or ill to which man's inventive craft may drive him one time or another is inside the human enclave and does not touch the nature of things.

The immunity of the whole, untroubled in its depth by the importunities of man, that is, the essential immutability of Nature as the cosmic order, was indeed the backdrop to all of mortal man's enterprises, between the abiding and the changing: the abiding was Nature, the changing his own works. The greatest of these works was the city, and on it he could offer some measure of abidingness by the laws he made for it and undertook to honor. But no long-range certainty pertained to this contrived abidingness. As a precarious artifact, it can lapse or go astray. Not even within its artificial space, with all the freedom it gives to man's determination of self, can the arbitrary ever supersede the basic terms of his being. The very inconstancy of human fortunes assures the constancy of the human condition. Chance and luck and folly, the great equalizers in human affairs, act like an entropy of sorts and make all definite designs in the long run revert to the perennial norm. Cities rise and fall, rules come and go, families prosper and decline; no change is there to stay, and in the end, with all the temporary deflections balancing each other out, the state of man is as it always was. So here too, in his very own artifact, man's control is small and his abiding nature prevails.

Still, in this citadel of his own making, clearly set off from the rest of things and entrusted to him, was the whole and sole domain of man's responsible action. Nature was not an object of human responsibility—she taking care of herself and, with some coaxing and worrying, also of man: not ethics, only cleverness applied to her. But in the city, where men deal with men, cleverness must be wedded to morality, for this is the soul of its being. In this intra-human frame dwells all traditional ethics and matches the nature of action delimited by this frame.

II

Let us extract from the preceding those characteristics of human action which are relevant for a comparison with the state of things today.

1. All dealing with the non human world, i.e., the whole realm of *technē* (with the exception of medicine), was ethically neutral—in respect both of the object and the subject of such action: in respect of the object, because it impinged but little on the self-sustaining nature of things and thus raised no question of permanent injury to the integrity of its object, the natural order as a whole; and in respect of the subject it was ethically neutral because *technē* as an activity conceived itself as a determinate tribute to necessity and not as an indefinite, self-validating advance to mankind's major goal, claiming in its pursuit man's ultimate effort and concern. The real vocation of man lay elsewhere. In brief, action on nonhuman things did not constitute a sphere of authentic ethical significance.

2. Ethical significance belonged to the direct dealing of man with man, including the dealing with himself: all traditional ethics is *anthropocentric*.

3. For action in this domain, the entity "man" and his basic condition was considered constant in essence and not itself an object of reshaping *technē*.

4. The good and evil about which action had to care lay close to the act, either in the praxis

itself or in its immediate reach, and were not a matter for remote planning. This proximity of ends pertained to time as well as space. The effective range of action was small, the time-span of foresight, goal-setting, and accountability was short, control of circumstances limited. Proper conduct had its immediate criteria and almost immediate consummation. The long run of consequences beyond was left to change, fate, or providence. Ethics accordingly was of the here and now, of occasions as they arise between men, of the recurrent, typical situations of private and public life. The good man was he who met these contingencies with virtue and wisdom, cultivating these powers in himself, and for the rest resigning himself to the unknown.

All enjoinders and maxims of traditional ethics, materially different as they may be, show this confinement to the immediate setting of the action. "Love thy neighbor as thyself"; "Do unto others as you would wish them to do unto you"; "Instruct your child in the way of truth"; "Strive for excellence by developing and actualizing the best potentialities of your being *qua* man"; "Subordinate your individual good to the common good"; "Never treat your fellow man as a means only but always *also* as an end in himself"—and so on. Note that in all those maxims the agent and the "other" of his action are sharers of a common present. It is those alive now and in some commerce with me that have a claim on my conduct as it affects them by deed or omission. The ethical universe is composed of contemporaries, and its horizon to the future is confined by the foreseeable span of their lives. Similarly confined is its horizon of place, within which the agent and the other meet as neighbor, friend or foe, as superior and subordinate, weaker and stronger, and in all the other roles in which humans interact with one another. To this proximate range of action all morality was geared.

III

It follows that the *knowledge* that is required—besides the moral will—to assure the morality of action fitted these limited terms; it was not the knowledge of the scientist or the expert, but knowledge of a kind readily available to all men of good will. Kant went so far as to say that "human reason can, in matters of morality, be easily brought to a high degree of accuracy and completeness even in the most ordinary intelligence";[1] that "there is no need of science or philosophy for knowing what man has to do in order to be honest and good, and indeed to be wise and virtuous.... [Ordinary intelligence] can have as good a hope of hitting the mark as any philosopher can promise himself";[2] and again: "I need no elaborate acuteness to find out what I have to do so that my willing be morally good. Inexperienced regarding the course of the world, unable to anticipate all the contingencies that happen in it," I can yet know how to act in accordance with the moral law.[3]

Not every thinker in ethics, it is true, went so far in discounting the cognitive side of moral action. But even when it received much greater emphasis, as in Aristotle, where the discernment of the situation and what is fitting for it makes considerable demands on experience and judgment, such knowledge has nothing to do with the science of things. It implies, of course, a general conception of the human good as such, a conception predicated on the presumed invariables of man's nature and condition, which may or may not find expression in a theory of its own. But its translation into practice requires a knowledge of the here and now, and this is entirely nontheoretical. This "knowledge" proper to virtue (of the "where, when, to, whom, and how") stays with the immediate issue, in whose defined context the action *as the agent's own* takes its course and within which it terminates. The good or bad of the action is wholly decided within that short-term context. Its moral quality shines forth from it, visible to its witnesses. No one was held responsible for the unintended later affects of his well-intentioned, well-considered, and well-performed act. The short arm of human power did not call for a long arm of predictive knowledge; the shortness of the one is as little culpable as that of the other. Precisely because the human good, known in its generality, is the same for all

time, its relation or violation takes place at each time, and its complete locus is always the present.

IV

All this has decisively changed. Modern technology has introduced actions of such novel scale, objects, and consequences that the framework of former ethics can no longer contain them. The *Antigone* chorus on the *deinotes*, the wondrous power, of man would have to read differently now; and its admonition to the individual to honor the laws of the land would no longer be enough. To be sure, the old prescriptions of the "neighbor" ethics—of justice, charity, honesty, and so on—still hold in their intimate immediacy of the nearest, day by day sphere of human interaction. But this sphere is overshadowed by a growing realm of collective action where doer, deed, and effect are no longer the same as they were in the proximate sphere, and which by the enormity of its powers forces upon ethics a new dimension of responsibility never dreamt of before.

Take, for instance, as the first major change in the inherited picture, the critical *vulnerability* of nature to man's technological intervention—unsuspected before it began to show itself in damage already done. This discovery, whose shock led to the concept and nascent science of ecology, alters the very concept of ourselves as a causal agency in the larger scheme of things. It brings to light, through the effects, that the nature of human action has *de facto* changed, and that an object of an entirely new order—no less than the whole biosphere of the planet—has been added to what we must be responsible for because of our power over it. And of what surpassing importance an object, dwarfing all previous objects of active man! Nature as a human responsibility is surely a *novum* to be pondered in ethical theory. What kind of obligation is operative in it? Is it more than a utilitarian concern? Is it just prudence that bids us not to kill the goose that lays the golden eggs, or saw off the branch on which we sit? But the "we" that here sits and may fall into the abyss is all future mankind, and the survival of the species is more than a prudential duty of its present members. Insofar as it is the fate of *man*, as affected by the condition of nature, which makes us care about the preservation of nature, such care admittedly still retains the anthropocentric focus of all classical ethics. Even so, the difference is great. The containment of nearness and contemporaneity is gone, swept away by the spatial spread and time-span of the cause-effect trains which technological practice sets afoot, even when undertaken for proximate ends. Their irreversibility conjoined to their aggregate magnitude injects another novel factor into the moral equation. To this take their cumulative character: their effects add themselves to one another, and the situation for later acting and being becomes increasingly different from what it was for the initial agent.

The cumulative self-propagation of the technological change of the world thus constantly overtakes the conditions of its contributing acts and moves through none but unprecedented situations, for which the lessons of experience are powerless. And not even content with changing its beginning to the point of unrecognizability, the cumulation as such may consume the basis of the whole series, the very condition of itself. All this would have to be co-intended in the will of the single action if this is to be a morally responsible one. Ignorance no longer provides it with an alibi.

Knowledge, under these circumstances, becomes a prime duty beyond anything claimed for it heretofore, and the knowledge must be commensurate with the causal scale of our action. The fact that it cannot really be thus commensurate, i.e., that the predictive knowledge falls behind the technical knowledge which nourishes our power to act, itself assumes ethical importance. Recognition of ignorance becomes the obverse of the duty to know and thus part of the ethics which must govern the ever more necessary self-policing of our out-sized might. No previous ethics had to consider the global condition of human life and the far-off future, even existence, of the race. There now being an issue demands, in brief, a new concept of duties and rights, for which previous ethics and metaphysics provide not even the principles, let alone a ready doctrine.

And what if the new kind of human action would mean that more than the interest of man alone is to be considered—that our duty extends farther and the anthropocentric confinement of former ethics no longer holds? It is at least not senseless anymore to ask whether the condition of extra-human nature, the biosphere as a whole and in its parts, now subject to our power, has become a human trust and has something of a more claim on us not only for our ulterior sake but for its own and in its own right. If this were the case it would require quite some rethinking in basic principles of ethics. It would mean to seek not only the human good, but also the good of things extra-human, that is, to extend the recognition of "ends in themselves" beyond the sphere of man and make the human good include the care for them. For such a role of stewardship no previous ethics has prepared us—and the dominant, scientific view of *Nature* even less. Indeed, the latter emphatically denies us all conceptual means to think of Nature as something to be honored, having reduced it to the indifference of necessity and accident, and divested it of any dignity of ends. But still, a silent plea for sparing its integrity seems to issue from the threatened plenitude of the living world. Should we heed this plea, should we grant its claim as sanctioned by the nature of things, or dismiss it as a mere sentiment on our part, which we may indulge as fair as we wish and can afford to do? If the former, it would (if taken seriously in its theoretical implications) push the necessary rethinking beyond the doctrine of action, i.e., ethics, into the doctrine of being, i.e., metaphysics, in which all ethics must ultimately be grounded. On this speculative subject I will here say no more than that we should keep ourselves open to the thought that natural science may not tell the whole story about Nature.

V

Returning to strictly intra-human considerations, there is another ethical aspect to the growth of *techne* as a pursuit beyond the pragmatically limited terms of former times. Then, so we found, *techne* was a measured tribute to necessity, not the road to mankind's chosen goal—a means with a finite measure of adequacy to well-defined proximate ends. Now, *techne* in the form of modern technology has turned into an infinite forward-thrust of the race, its most significant enterprise, in whose permanent, self-transcending advance to ever greater things the vocation of man tends to be seen, and whose success of maximal control over things and himself appears as the consummation of his destiny. Thus the triumph of *Homo faber* over his external object means also his triumph in the internal constitution of *Homo sapiens*, of whom he used to be a subsidiary part. In other words, technology, apart from its objective works, assumes ethical significance by the central place it now occupies in human purpose. Its cumulative creation, the expanding artificial environment, continuously reinforces the particular powers in man that created it, by compelling their unceasing inventive employment in its management and further advance, and by rewarding them with additional success—which only adds to the relentless claim. This positive feedback of functional necessity and reward—in whose dynamics pride of achievement must not be forgotten—assures the growing ascendancy of one side of man's nature over all the others, and inevitably at their expense. If nothing succeeds like success, nothing also entraps like success. Outshining in prestige and starving in resources whatever else belongs to the fullness of man, the expansion of his power is accompanied by a contraction of his self-conception and being. In the image he entertains of himself—the potent self-formula which determines his actual being as much as it reflects it—man now is evermore the maker of what he has made and the doer of what he can do, and most of all the preparer of what he will be able to do next. But not you or I: it is the aggregate, not the individual doer or deed that matters here; and the indefinite future, rather than the contemporary context of the action, constitutes the relevant horizon of responsibility. This requires imperatives of a new sort. If the realm of making has invaded the space of essential action, then morality must invade the realm of making, from which it had formerly stayed aloof, and must do so in the form of public policy. With issues of such inclusiveness and such lengths of

anticipation public policy has never had to deal before. In fact, the changed nature of human action changes the very nature of politics.

For the boundary between "city" and "nature" has been obliterated: the city of men, once an enclave in the non human world, spreads over the whole of terrestrial nature and usurps its place. The difference between the artificial and the natural has vanished, the natural is swallowed up in the sphere of the artificial, and at the same time the total artifact, the works of man working on and through himself, generates a "nature" of its own, i.e., a necessity with which human freedom has to cope in an entirely new sense. Once it could be said *Fiat justitia, pereat mundus*, "Let justice be done, and may the world perish"—where "world," of course, meant the renewable enclave in the imperishable whole. Not even rhetorically can the like be said anymore when the perishing of the whole through the doings of man—be they just or unjust—has become a real possibility. Issues never legislated on come into the purview of the laws which the total city must give itself so that there will be a world for the generations of man to come.

That there *ought to* be through all future time such a world fit for human habitation, and that it ought in all future time to be inhabited by a mankind worthy of the human name, will be readily affirmed as a general axiom or a persuasive desirability of speculative imagination (as persuasive and undemonstrable as the proposition that there being a world at all is "better" than there being none): but as a *moral* proposition, namely, a practical *obligation* toward the posterity of a distant future, and a principle of decision in present action, it is quite different from the imperatives of the previous ethics of contemporaneity; and it has entered the moral scene only with our novel powers and range of prescience.

The *presence of man in the world* had been a first and unquestionable given, from which all idea of obligation in human conduct started out. Now it has itself become an *object* of obligation—the obligation namely to ensure the very premise of all obligation, i.e., the *foothold* for a moral universe in the physical world—the existence of mere *candidates* for a moral order. The difference this makes for ethics may be illustrated in one example.

VI

Kant's categorical imperative said: "Act so that you *can* will that the maxim of our action be made the principle of a universal law." The "can" here invoked is that of reason and its consistency with itself: *Given* the existence of a community of human agents (acting rational beings), the action must be such that it can without self-contradiction be imagined as a general practice of that community. Mark that the basic reflection of morals here is not itself a moral but a logical one: The "I *can* will" or "I *cannot* will" expresses logical compatibility or incompatibility, not moral approbation or revulsion. But there is no self-contradiction in the thought that humanity would once come to an end, therefore also none in the thought that the happiness of present and proximate generations would be bought with the unhappiness or even non existence of later ones—as little as, after all, in the inverse thought that the existence or happiness of later generations would be bought with the unhappiness or even partial extinction of present ones. The sacrifice of the future for the present is *logically* no more open to attack than the sacrifice of the present for the future. The difference is only that in the one case the series goes on, and in the other it does not. But that it *ought to go on*, regardless of the distribution of happiness or unhappiness, even with a persistent preponderance of unhappiness over happiness, nay, even of immorality over morality[8]—this cannot be derived from the rule of self-consistency *within* the series, long or short as it happens to be: it is a commandment of a very different kind, lying outside and "prior" to the series as a whole, and its ultimate grounding can only be metaphysical.

An imperative responding to the new type of human action and addressed to the new type of agency that operates it might run thus: "Act so that the effects of your action are compatible with the permanence of genuine human life"; or expressed negatively: "Act so that the effects of

your action are not destructive of the future possibility of such life"; or simply: "Do not compromise the conditions for an indefinite continuation of humanity on earth"; or most generally: "In your present choices, include the future wholeness of Man among the objects of your will."

It is immediately obvious that no rational contradiction is involved in the violation of this kind of imperative. I *can* will the present good with sacrifice of the future good. It is also evident that the new imperative addresses itself to public policy rather than private conduct, which is not in the causal dimension to which that imperative applies. Kant's categorical imperative was addressed to the individual, and its criterion was instantaneous. It enjoined each of us to consider what would happen *if* the *maxim* of my present action were made, or at this moment already were, the principle of a universal legislation; the self-consistency or inconsistency of such a *hypothetical* universalization is made the test for my *private* choice. But it was no part of the reasoning that there is any probability of my private choice *in fact* becoming universal law, or that it might contribute to its becoming that. The universalization is a thought-experiment by the private agent not to test the immanent morality of his action. Indeed, real consequences are not considered at all, and the principle is one not of objective responsibility but of the subjective quality of my self-determination. The new imperative invokes a different consistency: not that of the act with itself, but that of its eventual *effects* with the continuance of human agency in times to come. And the "universalization" it contemplates is by no means hypothetical—i.e., a purely logical transference from the individual "me" to an imaginary, causally unrelated "all" ("*if* everybody acted like that"); on the contrary, the actions subject to the new imperative—actions of the collective whole—have their universal reference in their actual scope of efficacy: they "totalize" themselves in the progress of their momentum and thus are bound to terminate in shaping the universal dispensation of things. This adds a *time* horizon to the moral calculus which is entirely absent from the instantaneous logical operation of the Kantian imperative: whereas the latter extrapolates into an ever-present order of abstract compatibility, our imperative extrapolates into a predictable real *future* as the open-ended dimension of our responsibility.

VII

Similar comparisons could be made with all the other historical forms of the ethics of contemporaneity and immediacy. The new order of human action requires a commensurate ethics of foresight and responsibility, which is as new as are the issues with which it has to deal. We have seen that these are the issues posed by the works of *Homo faber* in the age of technology. But among those novel works we haven't mentioned yet the potentially most ominous class. We have considered *techne* only as applied to the non-human realm. But man himself has been added to the objects of technology. *Homo faber* is turning upon himself and gets ready to make over the maker of all the rest. This consummation of his power, which may well portend the overpowering of man, this final imposition of art on nature, calls upon the utter resources of ethical thought, which never before has been faced with elective alternatives to what were considered the definite terms of the human condition.

a. Take, for instance, the most basic of these "givens," man's mortality. Who ever before had to make up his mind on its desirable and *eligible* measure? There was nothing to choose about the upper limit, the "three score years and ten, or by reason of strength fourscore." Its inexorable rule was the subject of lament, submission, or vain (not to say foolish) wish-dreams about possible exceptions—strangely enough, almost never of affirmation. The intellectual imagination of a George Bernard Shaw and a Jonathan Swift speculated on the privilege of not having to die, or the curse of not being able to die. (Swift with the latter was the more perspicacious of the two.) Myth and legend toyed with such themes against the acknowledged background of the unalterable, which made the earnest man rather pray "teach us to number our days that we may get a heart of wisdom" (Psalm 90). Nothing of this was in the realm of doing, and effective decision. The question was only how to relate to the stubborn fact.

But lately, the dark cloud of inevitability seems to lift. A practical hope is held out by certain advances in cell biology to prolong, perhaps indefinitely extend the span of life by counteracting biochemical processes of aging. Death no longer appears as a necessity belonging to the nature of life, but as an avoidable, at least in principle tractable and long-delayable, organic malfunction. A perennial yearning of mortal man seems to come nearer fulfillment. And for the first time we have in earnest to ask the question "How desirable is this? How desirable for the individual, and how for the species?" These questions involve the very meaning of our finitude, the attitude toward death, and the general biological significance of the balance of death and procreation. Even prior to such ultimate questions are the more pragmatic ones of who should be eligible for the boon: persons of particular quality and merit? of social eminence? those that can pay for it? everybody? The last would seem the only just course. But it would have to be paid for at the opposite end, at the source. For clearly, on a population-wide scale, the price of extended age must be a proportional slowing of replacement, i.e., a diminished access of new life. The result would be a decreasing proportion of youth in an increasingly aged population. How good or bad would that be for the general condition of man? Would the species gain or lose? And how *right* would it be to preempt the place of youth? Having to die is bound up with having been born: mortality is but the other side of the perennial spring of "a natality" (to use Hannah Arendt's term). This had always been ordained; now its meaning has to be pondered in the sphere of decision.

To take the extreme (not that it will ever be obtained): if we abolish death, we must abolish procreation as well, for the latter is life's answer to the former, and so we would have a world of old age with no youth, and of known individuals with no surprises of such that had never been before. But this perhaps is precisely the wisdom in the harsh dispensation of our mortality: that it grants us the eternally renewed promise of the freshness, immediacy, and eagerness of youth, together with the supply of otherness as such. There is no substitute for this in the greater accumulation of prolonged experience: it can never recapture the unique privilege of seeing the world for the first time and with new eyes, never relive the wonder which, according to Plato, is the beginning of philosophy, never the curiosity of the child, which rarely enough lives on as thirst for knowledge in the adult, until it wanes there too. This ever renewed beginning, which is only to be had at the price of ever repeated ending, may well be mankind's hope, its safeguard against lapsing into boredom and routine, its chance of retaining the spontaneity of life. Also, the role of the *memento mori* in the individual's life must be considered, and what its attenuation to indefiniteness may do to it. Perhaps a non-negotiable limit to our expected time is necessary for each of us as the incentive to number our days and make them count.

So it could be that what by intent is a philanthropic gift of science to man, the partial granting of his oldest wish—to escape the curse of mortality—turns out to be to the detriment of man. I am not indulging in prediction and, in spite of my noticeable bias, not even in valuation. My point is that already the promised gift raises questions that had never to be asked before in terms of practical choice, and that no principle of former ethics, which took the human constants for granted, is competent to deal with them. And yet they must be dealt with ethically and by principle and not merely by the pressure of interest.

b. It is similar with all the other, quasi-utopian powers about to be made available by the advances of biomedical science as they are translated into technology. Of these, *behavior control* is much nearer to practical readiness than the still hypothetical prospect I have just been discussing, and the ethical questions it raises are less profound but have a more direct bearing on the moral conception of man. Here again, the new kind of intervention exceeds the old ethical categories. They have not equipped us to rule, for example, on mental control by chemical means or by direct electrical action of the brain via implanted electrodes—undertaken, let us assume, for defensible and even laudable ends. The mixture of beneficial and dangerous potentials is obvious, but the lines are not easy to draw. Relief of mental patients from

distressing and disabling symptoms seems unequivocally beneficial. But from the relief of the *patient*, a goal entirely in the tradition of the medical art, there is an easy passage to the relief of *society* from the inconvenience of difficult individual behavior among its members: that is, the passage from medical to social application; and this opens up an indefinite field with grave potentials. The troublesome problems of rule and unruliness in modern mass society make the extension of such control methods to non-medical categories extremely tempting for social management. Numerous questions of human rights and dignity arise. The difficult question of preemption care versus enabling care insists on concrete answers. Shall we induce learning attitudes in school children by the mass administration of drugs, circumventing the appeal to autonomous motivation? Shall we overcome aggression by electronic pacification of brain areas? Shall we generate sensations of happiness or pleasure or at least contentment through independent stimulation (or tranquilizing) of the appropriate centers—independent, that is, of the objects of happiness, pleasure, or content and their attainment in personal living and achieving? Candidacies could be multiplied. Business firms might become interested in some of these techniques for performance-increase among their employees.

Regardless of the question of compulsion or consent, and regardless also of the question of undesirable side-effects, each time we thus bypass the human way of dealing with human problems, short-circuiting it by an impersonal mechanism, we have taken away something from the dignity of personal selfhood and advanced a further step on the road from responsible subjects to programmed behavior systems. Social functionalism, important as it is, is only one side of the question. Decisive is the question of what kind of individuals the society is composed of to make its existence valuable as a whole. Somewhere along the line of increasing social manageability at the price of individual autonomy, the question of the worthwhileness of the human enterprise must pose itself. Answering it involves the image of man we entertain. We must think it anew in light of the things we can do to it now and could never do before.

c. This holds even more with respect to the last object of a technology applied on man himself—the genetic control of future men. This is too wide a subject for cursory treatment. Here I merely point to this most ambitious dream of *Homo faber*, summed up in the phrase that man will take his own evolution in hand, with the aim of not just preserving the integrity of the species but of modifying it by improvements of his own design. Whether we have the right to do it, whether we are qualified for that creative role, is the most serious question that can be posed to man finding himself suddenly in possession of such failed powers. Who will be the imagemakers, by what standards, and on the basis of what knowledge? Also, the question of the moral right to experiment on future human beings must be asked. These and similar questions, which demand an answer before we embark on a journey into the unknown, show most vividly how far our powers to act are pushing us beyond the terms of all former ethics.

VIII

The ethically relevant common feature in all the examples adduced is what I like to call the inherently "utopian" drift of our actions under the conditions of modern technology, whether it works on nonhuman or on human nature, and whether the "utopia" at the end of the road be planned or unplanned. By the kind and size of its snowballing effects, technological power propels us into goals of a type that was formerly the preserve of Utopias. To put it differently, technological power has turned what used and ought to be tentative, perhaps enlightening, plays of speculative reason into competing blueprints for projects, and in choosing between them we have to choose between extremes of remote effects. The one thing we can really know of them is their extremism as such—that they concern the total condition of nature on our globe and the very kind of creatures that shall, or shall not, populate it. In consequence of the inevitably "utopian" scale of modern technology, the salutary gap between everyday and ultimate issues, between occasions, is closing. Living now constantly in the shadow of

unwanted, built-in, automatic utopianism, we are constantly confronted with issues whose positive choice requires supreme wisdom—an impossible, and in particular for contemporary man, who denies the very existence of its object: viz., objective value and truth. We need wisdom most when we believe in it least. If the new nature of our acting then calls for a new ethics of long-range responsibility, coextensive with the range of our power, it calls in the name of that very responsibility also for a new kind of humility—a humility not like former humility, i.e., owing to the littleness, but owing to the excessive magnitude of our power, which is the excess of our power to act over our power to foresee and our power to evaluate and to judge. In the face of the quasieschatological potentials of our technological processes, ignorance of the ultimate implications becomes itself a reason for responsible restraint—as the second best to the possession of wisdom itself.

One other aspect of the required new ethics of responsibility for and to a distant future is worth mentioning: the insufficiency of representative government to meet the new demands on its normal principles and by its normal mechanics. For according to these, only *present* interests make themselves heard and felt and enforce their condition. It is to them that public agencies are accountable, and this is the way in which concretely the respecting of rights comes about (as distinct from their abstract acknowledgment). But the *future* is not represented, it is not a force that can throw its weight into the scales. The nonexistent has no lobby, and the unborn are powerless. Thus accountability to them has no political reality behind it yet in present decision-making, and when they can make their complaint, then we, the culprits, will no longer be there.

This raises to an ultimate pitch the old question of the power of the wise, or the force of ideas not allied to self-interest, in the body politic. What *force* shall represent the future in the present? However, before *this* question can become earnest in practical terms, the new ethics must find its theory, on which dos and don'ts can be based. the question of what *insight* or value-knowledge shall represent the future in the present.

IX

And here is where I get stuck, and where we all get stuck. For the very same movement which put us in possession of the powers that have now to be regulated by norms—the movement of modern knowledge called science—has by a necessary complementarity eroded the foundations from which norms could be derived; it has destroyed the very idea of norm as such. Not, fortunately, the feeling for norm and even for particular norms. But this feeling became uncertain of itself when contradicted by alleged knowledge or at least denied all sanction by it. Anyway and always does it have a difficult enough time against the loud clamors of greed and fear. Now it must in addition blush before the frown of superior knowledge, as unfounded and incapable of foundation. First, Nature has been "neutralized" with respect to value, then man himself. Now we shiver in the nakedness of a nihilism in which near-omnipotence is paired with near-emptiness, greatest capacity with knowing least what for. With the apocalyptic pregnancy of our actions, that very knowledge which we lack has become more urgently needed than at any other stage in the adventure of mankind. Alas, urgency is no promise of success. On the contrary, it must be avowed that to seek for wisdom today requires a good measure of unwisdom. The very nature of the age which cries out for an ethical theory makes it suspiciously look like a fool's errand. Yet we have no choice in the matter but to try.

It is a question whether without restoring the category of the sacred, the category most thoroughly destroyed by the scientific enlightenment, we can have an ethics able to cope with the extreme powers which we possess today and constantly increase and are almost compelled to use. Regarding those consequences imminent enough still to hit ourselves, fear can do the job—so often the best substitute for genuine virtue or wisdom. But this means fails us toward the more distant prospects, which here matter the most, especially

smallness. Only awe of the sacred with its unqualified veto is independent to fit computations of mundane fear and the solace of uncertainty about distant consequences. But religion as a soul-determining force is no longer there to be summoned to the aid of ethics. The latter must stand on its worldly feet—that is, on reason and its fitness for philosophy. And while of faith it can be said that it either is there or is not, of ethics it holds that it must be there.

It must be there because men act, and ethics is for the reordering of actions and for regulating the power to act. It must be there all the more, then, the greater the powers of acting that are to be regulated; and with their size, the ordering principle must also fit their kind. Thus, novel powers to act require novel ethical rules and perhaps even a new ethics.

"Thou shalt not kill" was enunciated because man has the power to kill and often the occasion and even inclination for it—in short, because killing is actually done. It is only under the *pressure* of real habits of action, and generally of the fact that always action already takes place, without *this* having to be commanded first, that ethics as the ruling of such acting under the standard of the good or the permitted enters the stage. Such a *pressure* emanates from the novel technological powers of man, whose exercise is given with their existence. *If* they really are as novel in kind as here contended, and if by the kind of their potential consequences they really have abolished the moral neutrality which the technical commerce with matter hitherto enjoyed—then their pressure bids to seek for new prescriptions in ethics which are competent to assume their guidance, but which first of all can hold their own theoretically against that very pressure. To the demonstration of those premises this paper was devoted. If they are accepted, then we who make thinking our business have a task to last us for our time. We must do it in time, for since we act anyway we shall have some ethic or other in any case, and without a supreme effort to determine the right one, we may be left with a wrong one by default.

NOTES

1. Immanuel Kant, *Groundwork of the Metaphysics of Morals*, preface.
2. *Op. cit.*, chapter 1.
3. *Ibid.* (I have followed H. J. Paton's translation with some changes.)
4. On this last point, the biblical God changed his mind to an all-encompassing "yes" after the Flood.

1.3.2 CONSTRUCTING THE PROFESSIONAL RESPONSIBILITY OF ENGINEERS

MICHAEL DAVIS

The previous selection by Hans Jonas (1.3.1) emphasized the moral imperative for humanity to "take responsibility" for the effects of technology on the nonhuman world and on ourselves, and to develop a new form of ethics that encompasses a "long-range responsibility" commensurate with the awesome powers that technology affords us. But Jonas does not tell us how to understand the concept of responsibility, or how this concept can be applied to the work of engineers and other technology professionals.

Responsibility is a term that has several distinct meanings that Davis teases apart in this article. He distinguishes several types of causal responsibility from

Source: From "Science Engineering Ethics", Volume 18, 2012, Michael Davis, "Ain't No One Here But Us Social Forces": *Constructing the Professional Responsibility of Engineers*, pp. 13–34. Copyright © Springer Science+Business Media B.V. 2010. Reprinted with kind permission from Springer Science+Business Media B.V.

"responsibility-as-liability" for praise or blame, from "responsibility-as-competency," and from "responsibility-as-office," which in turn is distinct from "responsibility-as-a-domain-of-tasks." Because of these different senses of responsibility, discussions about who is responsible for what are often confusing, particularly in organizations where many individuals contribute in a partial way to producing some outcome or bringing about some state of affairs. But Davis argues that none of the usual arguments that are used to avoid attributions of responsibility to engineers are sound, and that in fact, there is always the option for individuals and groups to "take responsibility" for the effects that their work has on the world. While somewhat challenging, this article repays the effort of reading it carefully.

FOCUS QUESTIONS

1. What is the difference between the backward-looking and the forward-looking senses of responsibility? Provide some examples of each kind.
2. What are the four necessary conditions of responsibility-as-liability that Davis identifies? What reasons does he give for thinking these criteria are not really necessary for ascriptions of moral responsibility?
3. What are the seven arguments that are commonly used to defend the claim that engineers cannot be morally responsible for the effects of their work? Do any of them hold water in Davis's opinion?
4. How does Davis's analysis of the THERAC-25 disaster illustrate the argument that many causes can undercut ascriptions of moral responsibility?
5. What faults does Davis find with the argument from institutional constraints? What does the NSPE *Code of Ethics for Engineers* (Appendix B) say about serving the public interest as well as their employers and the acceptance of personal responsibility?
6. What does Davis say about the idea that technology is itself "value neutral," and that engineers cannot be responsible for all of the possible misuses of the products they design? Compare his analysis to that found in Andrew Feenberg's "Democratic Rationalization" (Selection 1.2.2).
7. What reasons does Davis give, in the end, for why engineers should rationally choose to take responsibility in cases where others might wish to avoid it? Do you find these arguments persuasive?

KEYWORDS

accountability, causation, engineering, ethics, liability, responsibility, THERAC-25

There are many ways to avoid responsibility. Among them are attributing one's act or its consequences to society, market, culture, God, the devil, or drink. For engineers, "technology" or "the organization" might seem to serve quite well. I mock this response as, "Ain't no one here but us social forces"-recalling the old Vaudeville joke. Of course, the joke depends on someone being there in addition to the chickens, someone who can properly take responsibility but tries to avoid it-by a device too desperate to succeed. The response thus mocked is not a joke. It presupposes that certain facts prevent assignment of responsibility. For example, some have claimed that engineers cannot be responsible for some consequences of their work because the technology is so complex that they cannot reasonably be expected to foresee those consequences.[1] A fact, the impossibility of reasonable foresight, rules out the responsibility in question.

What I propose to do here is argue that any such appeal to mere fact fails. The responsibility in question is something engineers can, and generally do, take on by their own voluntary acts.

Engineers construct their responsibility, much of it at least, rather than discover it or have it imposed on them. Taking responsibility is an important but, it seems, overlooked way of becoming responsible. For engineers, the proper question is (primarily): *what responsibility have* we *taken on*? The responsibility of engineers is, in large part, what, as a profession, they make it.

I shall have nothing to say here about whether engineers have taken on too much, too little, or just the right amount of responsibility. My focus is on arguments supposedly defeating responsibility claims.

PRELIMINARY CAVEATS

Like most important terms in ethics, "responsibility" has too many senses for the term's safe use without distinguishing the more important ones, understanding their relationship, and then trying to keep them straight. We may distinguish the following nine interrelated senses as relevant here (all of which I will eventually need):

1. responsibility-as-simple-causation ("Katrina was responsible for the flooding in New Orleans"),
2. responsibility-as-faulty-causation ("responsible because he acted carelessly"),
3. responsibility-as-good-causation ("She's responsible for our success, give her the credit"),[2]
4. responsibility-as-competency ("a responsible person"—rather than, say, a careless or incompetent one—responsibility as a disposition or virtue),
5. responsibility-as-power ("your skill gives you a responsibility"),[3]
6. responsibility-as-office ("the person responsible", the one in charge),
7. responsibility-as-domain-of-tasks ("these are her responsibilities", that is, what she is supposed to do),
8. responsibility-as-liability ("you should *pay* because you are responsible for what happened"), and
9. responsibility-as-accountability ("you should *explain* because you are responsible for what happened").[4]

This list has more structure than may be apparent at first. Five of the nine senses are (more or less) factual: the simple-cause, competency, power, office, and task senses. Two, in addition to the simple-cause sense, are entirely historical (or "backward looking"): the faulty-cause and the good-cause senses. When we are talking about responsibility in any of the three historical senses, we should (and generally do) use the past tense: "She *was* responsible for that." The remaining two (the liability and accountability senses) seem to be about the future, that is, they point to what should be done. In this respect, they resemble the power, office, and task senses. The responsibility that comes with power, office, or task is a responsibility to do something (to benefit humanity, to oversee this project, or to perform these tasks). The responsibility that comes with liability (whatever explains the liability) is also something to do, for example, accept blame, apologize, make compensation, or perform as promised. The responsibility that comes with accountability is to give an account of something, for example, offer a justification, excuse, or explanation, or answer pertinent questions. The standard way to talk about responsibility in one of these five future-oriented senses (power, office, task, liability, or accountability) is in the present tense: "She *is* responsible for that."

This difference in tense is only a general index. It is easy to find causal uses in the present tense (when there is a continuing or continual causing) as in, "He is responsible for all my troubles [that is, the cause of them all]." It is also easy to find uses of the future-oriented senses in the past tense (when the power has been lost, the office vacated, the task completed, the liability satisfied, or the account given) as in, "He *was* responsible for paying the damages [but paid them long ago]." Nonetheless, I think it important to stress the future-oriented aspect of all five senses, though each typically presupposes something in the past (the act assigning office or tasks, the event imposing liability or imparting power, and so on).[5] What is important about the future-oriented senses of responsibility is

precisely that they have something to say about the future, something not true of competency or any of the three historical senses.

Neither technology nor organization seems to threaten responsibility in any of the five factual senses (cause, competency, power, office, task). Whatever technology or organization becomes, engineers who help to design, build, maintain, or dispose of technology will remain causal agents, responsible in the first (simple-cause) sense of responsibility for whatever they bring about (whether anyone can trace their individual contributions or not). They may also remain responsible persons (agents to be trusted), experts (with the power that comes from knowledge, skill, or judgment), persons responsible (agents in charge), and persons with responsibilities (tasks to do). When technology or organization seems to threaten "responsibility", it is responsibility-as-liability or responsibility-as-accountability that is threatened. In particular, it might be (and sometimes is) argued that technology may allow engineers to do great harm without fault. We cannot (it is said) hold engineers liable or even accountable for the harm they do without fault because fault is a logical precondition of responsibility in the liability or accountability sense. What engineers do without fault is not their responsibility.[6]

Though strong, the relation between fault and liability (or accountability) is, I think, not that of *logical* precondition. The law is quite capable of holding persons responsible for outcomes when they were not personally at fault—without any incoherence. Among terms covering this sort of responsibility are "strict liability" (liability when there is causal agency but no fault), "vicarious liability" (liability when the fault is in some person other than the one held liable, such as an agent or employee), and "surety liability" (when liability depends on contract rather than fault, causal agency, or relation of agent to principal).[7]

Our question, though, is not what the law can do with responsibility-as-liability (much less with responsibility-as-accountability). Our question is in part what the law *should* do with responsibility—and, in part, what we should do with it outside the law, for example, when attributing professional responsibility. This "should" depend on both

There is a weak, intermediate, and strong sense both of "moral justification" and "rational justification". I can, for example, morally justify my conduct by showing (a) that morality does not forbid it (a weak justification), (b) by showing that morality recommends it (an intermediate justification), or (c) by showing that morality requires it (a strong justification). We need not worry about this detail here. Weak moral justification is all we need for the arguments in question. Only at the end of this paper will I have to say anything about rational justification.

While our question is (primarily) about moral responsibility rather than legal, these examples of legal responsibility-as-liability are worth keeping in mind. Since they seem not to stretch the concept of responsibility to the breaking point, or indeed to stretch it much at all, they demonstrate the range of that concept. Any claim that *moral* responsibility is not as flexible as legal responsibility will have to rely on moral arguments, not on appeal to the "logic of responsibility" (much less to "what 'responsibility' means"). I therefore reject John Ladd's claim that moral responsibility is separate from the nine senses just distinguished. Ladd defines moral responsibility as (roughly) "the concern of people for people."[9] Insofar as this concern constitutes responsibility at all, it seems to me to define a domain of (morally required or recommended) tasks (say, taking people into account when acting) which we also find in law. While I do not think moral responsibility and legal responsibility *necessarily* mirror each other, I see no reason to suppose that the term "responsibility" in morality has a special meaning that cannot be incorporated into law—or that "responsibility" in law has a special meaning not to be found in morality. The difference has certainly not been established.

There are certain factors commonly put forward as conditions *necessary* for the assignment of *moral* responsibility-as-liability—and of moral responsibility-as-accountability when distinguished. The four most common are: (1) cause (or causal factor), (2) rationality (or competency), (3) choice (or freedom), and (4) knowledge (or foresight). (See for example, Swierstra and Jelsma 2006, esp. 312–313.) A person who does not cause an event cannot (it is said) be morally responsible for it, nor can a person who is not a "responsible agent", nor can a person who had no

happen (or at least what she risked). Part of what makes this collection of factors attractive is that each connects responsibility-as-liability with another sense of responsibility. Thus, the first explains why we might use "responsibility" for both liability and cause (any liability presupposes cause); the second, our use of "responsible person" for someone who can be held liable (a competent agent); the third and fourth, with fault, since fault requires both choice (an exercise of deliberative power) and something the agent knew or should have known (an exercise of cognitive power).[10] Indeed, this way of thinking about responsibility puts liability at the concept's center, with the other senses of responsibility derivable by subtracting one or more conditions of liability.[11]

However satisfying this structuring of the concept is, it should not be accepted. All the claims on which it relies seem to be false or, at least, no more than roughly true. Consider the requirement of cause: it seems not to fit some *failures* to rescue. We sometimes hold a person morally responsible for the death of another (in the liability or accountability sense) when she could have, but knowingly did not, come to the aid of the other, for example, a drowning child she could have saved with little risk, cost, or trouble to herself and no breach of duty to another. We blame her for the death though it would have occurred even if she had not been there to help, indeed, even if she had died a decade before. Her inaction, though certainly an event, indeed, a conscious act, is no more a cause of the child's death (whether "cause" means the most important factor or just a necessary one) than the non-presence of the billions of other people who might have saved the child had they been there.[12] Her non-rescue of the child has no causal significance (no special role in the actual process leading to death). Only her rescue of the child would have had causal significance—and that did not occur. We nonetheless hold her morally (and, in some jurisdictions, legally) responsible for the child's death (in the liability or accountability sense). If responsibility-as-simple-cause matters here (as I think it does), it is not as cause in this world but as part of an explanation invoking possible worlds. The cause in question exists only in at least one counterfactual world (one reasonably close to the actual one). Such a counterfactual cause cannot be the actual cause of the child's death. What is important to the responsibility of the bystander who did not help is that she *could* have saved the child but did not, that is, not faulty-cause in this world (since she caused nothing) but good-cause in another.[13]

Rationality is also said to be a necessary condition for moral responsibility. That is why (it is said) we do not hold the insane, the very young, or dumb animals morally responsible or, at least, why we do not *properly* hold them morally responsible. Of course, sometimes we do hold them morally responsible. For example, we may criticize even a young child for breaking a promise (saying something like, "You should have kept your promise"). Why we do that is open to discussion. My point now is merely that we do criticize the very young in this way, criticizing them in this way is blaming them (as criticizing a dog is not), blaming someone for breaking a promise is (normally) moral blame, and imposing moral blame on the very young at least seems to be holding them morally responsible (in the responsibility-as-liability sense). Certainly, when my son was four and broke a promise, I spoke to him as I would to someone who was morally responsible—and he responded accordingly, for example, by offering an excuse, justification, or otherwise accounting for his conduct, or by offering an apology or suggesting some other way to satisfy the liability incurred. Though I held him responsible in this way, I did not regard him as a competent adult (a responsible person). At four, he was far from that. On matters of safety, especially, I had more faith in our dog (who was ten). If we in fact sometimes treat even the very young as morally responsible in this way without supposing them to meet the rationality condition (as I believe we do), the hard question is whether doing so is morally proper (and, if it is, why). I need not answer that question here (though it seems proper enough to me). My purpose is to raise doubts about these four supposedly necessary conditions for responsibility-as-liability (and accountability). I have, I think, now said enough to raise doubts about this second one. We seem to act contrary to that condition without any feeling of doing something odd. Whatever this second condition is, it is not common sense. We can go to the third, choice.

We cannot (it is said) have any moral responsibility for an outcome when we had no choice but to bring it about. Consider, then, a minor traffic accident. A cat runs out from between parked cars just as I am passing. I slam on the brakes as quickly as I can—but too late. I have not been drinking, missed sleep, or otherwise done anything to impair my reactions. I was not speeding, using a cell phone, or otherwise failing to exercise reasonable care. The accident was, we might say, just one of those things—quite beyond my control—and therefore "definitely not my fault". Should I then drive on as if I were in no way responsible for the cat's death? I should not. I should stop, check to see if there is anything I can do for the cat, and perhaps call the police. I should do all that not to be a nice guy, a Good Samaritan, but because I am responsible (in some sense) for the cat's death. If I do not feel that responsibility, an obligation to respond in certain ways, I am a morally worse person than one who does.[14] I seem to owe the cat, or its owner, something for what I have done.[15] What I owe seems like responsibility-as-liability. Of course, I do not owe what I would have if the cat's death had been my fault (say, the result of negligence). But I still owe something. I shall offer no defense of this leap from responsibility-as-simple-cause to responsibility-as-liability here. I merely point it out as a fact of moral life, one that at least seems to undercut the claim that choice is *required* for moral responsibility.[16]

The cat also provides a good example of the fourth condition for moral responsibility (knowledge) not being met. I could not know that the cat was there, between the two cars, much less that it would run out just as I was driving by. Of course, I did know, or at least should have known, that such accidents are possible. Cars are big, fast, and hard to stop; cats, not always as careful of traffic as they should be. Yet, the risk I ran, however knowingly, merely by driving, does not seem to be enough for moral fault. Such general knowledge of unlikely events may help to explain strict liability in law, but it is not the sort of knowledge commonly presented to justify moral responsibility-as-liability (or accountability). More specific knowledge is supposed to be required, at least knowing that there is a reasonable chance that the event will happen if I drive (and the availability of alternatives both reasonable and safer). My knowledge fell well short of that.

I do not claim that these four conditions are irrelevant to responsibility-as-liability (or as-accountability) but only that they are not *necessary* conditions. They seem to me to function more like considerations relevant to rejecting, imposing, or accepting responsibility (mere reasons rather than requirements). The more of those conditions that are met, the stronger the case for moral responsibility, all else equal. But context is important too. In the right context, even one factor, such as simple cause, can be sufficiently strong (more or less by itself) to decide the question (as perhaps it is in the cat case).[17]

We should keep all this in mind as we evaluate arguments for the claim that technology or organization is undermining the responsibility of engineers for much they do. In particular, we should be charitable in our interpretation of these arguments. We should interpret them, if at all possible, as relying on a less questionable claim than that any of these four conditions is necessary for responsibility (though many do in fact seem to rely on that questionable claim). They fail even when they receive more charity than they deserve.

ARGUMENTS AGAINST HOLDING ENGINEERS RESPONSIBLE FOR WHAT THEY DO

There are at least seven arguments used to defend the claim that technology or organization is making, has made, or might someday make it impossible for engineers to be morally responsible for (at least some of) what they do. We may call them the argument from: (1) many hands, (2) many causes, (3) replaceability, (4) institutional constraint, (5) individual helplessness (lack of coordination), (6) individual ignorance, and (7) neutrality of technology. Each of these arguments suffers from one or more flaws. The responsibility in question is liability, accountability, or both.

Many Hands

We seem to owe the name of the first argument, if not the argument itself, to Dennis Thompson. His summary of it is worth quoting as much for what

it does not say as for what it does say: "Because many different officials contribute in many different ways to decisions and policies of government, it is difficult even in principle to identify who is morally responsible for political outcomes" (Thompson 1987, 41). For Thompson, the argument from many hands is primarily about large organizations ("governments"), not about technology. It is an epistemic argument, one about a problem more information might clear up, not a metaphysical argument (one relying on the nature of things). In addition, the problem the argument relies on makes identifying those morally responsible for an outcome "difficult", not impossible—and no newer than determining whether Henry II was responsible for the murder of Thomas Becket. Indeed, Thompson devotes the chapter that this quotation opens to assigning responsibility to politicians and government officials who appeal to one version or another of the argument to avoid it. Thompson is not afraid to look into the henhouse.

For Thompson, as for most who discuss the argument from many hands, a crucial assumption is that responsibility is primarily about holding *others* responsible.[18]

There is, I think, a close analogy between this view of responsibility as primarily imposed and what is called "the external" perspective on law. The external perspective on law misses essential features of law that the internal perspective includes, especially its authority. See especially, Hart (1961, pp. 101–102). From the outside, responsibility may seem to be primarily about having to answer to others for what one does (what Kant would call "heteronomy"). From the inside, responsibility may seem to be primarily about being trustworthy, someone capable of answering for what he does, a certain (virtuous) disposition (what is recognizably Kant's "autonomy").

So, for example, Helen Nissenbaum (whose subject *is* primarily technology) notes, "Boards of directors, task forces, or committees issue joint decisions, and on the occasions where these decisions are not universally approved by all their members but are the result of majority vote, we are left with the further puzzle of how to attribute responsibility."[19] Of course, "we" are left with the puzzle because "we" are not party to the "joint decisions" in question, not to decide our own responsibility. Were "we" among the decision-makers, we would at least know our part in the decisions, for example, whether we voted for or against the measure in question. The problem of many hands (so stated) does not seem to be much of a problem for holding *oneself* responsible.[20] In any case, it is no problem for someone who takes on the responsibility in question. I can take responsibility for what happened even if I have no idea how I contributed to the ultimate decision (as, for example, President Truman did when he said, "The buck stops here").

Many Causes

For decision-makers themselves, the problem is (in part at least) many *causes* rather than many *hands*. There are two versions of the resulting argument against responsibility. One relies on the complexity of the causal story; the other, on the problem of identifying "the cause" (or "the chief causes") even when the causal story is relatively simple. Like the problem of many hands, the problem of complexity is epistemic, but the problem of identifying "the cause" is metaphysical (a problem about what is to count as "cause").

Nissenbaum's discussion of the THERAC-25 disaster (three patients dead and three more seriously burned by what should have been a relatively safe means of irradiating cancer cells) provides a good example of the complexity version of the argument from many causes:

> After many months of study and trial-and-error testing, the origin of the malfunction was traced not to a single source, but to numerous faults, which included at least two significant software coding errors ("bugs") and a faulty microswitch. The impact of these faults was exacerbated by the absence of hardware interlocks, obscure error messages, inadequate testing and quality assurance, exaggerated claims about the reliability of the system in AECL's safety analysis and, in at least two cases, negligence on the parts of the hospitals where treatment was administered.[21]

In its complexity, the THERAC-25 disaster differs little from most engineering disasters. Because engineers try to design considerable redundancy into their systems, a disaster typically requires

cause, but several factors together constitute "the chief causes". Responsibility-as-faulty-cause must then be divided among these—and, insofar as each factor is the work of a different agent, among several agents. But the resulting responsibility-as-liability (or as-accountability) is not like weight, a burden *necessarily* lessened when so divided. Some parts of responsibility are lessened by being divided, especially financial liability, but the rest, such as the blame deserved or the obligation to explain one's own part, may not be.

Nonetheless, the THERAC-25 seems to illustrate the argument that many causes can undercut responsibility. After all, how could an engineer involved in the disaster honestly take any responsibility for the disaster (such as blame or the obligation to explain) until a post-disaster analysis identified the chief causes—and, if the causal story is too complex to identify any engineer's part, how could an engineer take any responsibility at all?

This question is not rhetorical. If it seems rhetorical, that is (in part at least) because we are not being careful enough about the sense of "responsibility" in question. If "responsibility" means responsibility-as-simple-cause or faulty-cause, then we cannot honestly take responsibility (that is, accept a description in which we are the cause) until we know the causal story and how our faulty conduct, if any, fits into it. The same seems true if we are to be called to account for the accident in question or held liable for it. How can an agent give an account or be held liable (even by himself) without knowing who was at fault?

The question is not rhetorical because both responsibility-as-liability and responsibility-as-accountability include more than responsibility for faulty conduct. Remember the moral responsibility arising simply from having caused the cat's death. We can see something like that in the THERAC-25 story. Confronted with the three deaths and three serious injuries, AECL (the manufacturer) investigated the THERAC-25's malfunction, eventually issuing a detailed post-disaster analysis. Undertaking that investigation meant accepting no more than that the THERAC-25 at least *seemed* to be a causal factor in the disaster. Undertaking such an investigation is something AECL's engineers should have urged. For engineers, even such a weak connection—being involved in the disaster as a possible causal factor—entails an obligation to prepare an accurate account of what happened. Preparing such an account is already taking (some) responsibility—and taking that responsibility (responsibility-as-accountability) is possible before anyone knows who was at fault—or even who caused the deaths and injuries.

This may seem like a small point, especially to engineers. What else would we do when there has been a disaster and we, or our organization, may be involved in it? Again, this is not a rhetorical question. For many in management or politics, the answer seems to be: *Call in the "spin doctors".* Even among professionals, there seem to be answers other than the one engineers take for granted. Lawyers, for example, generally do not undertake anything like AECL's post-disaster analysis when they lose a major case. They "move on", each lawyer taking whatever lesson she may from the experience.[22]

Whatever is true of other professionals, engineers consider it their responsibility to study any disaster that seems to arise from what they did—and to report what they find. To commit a certain mistake once, even a serious one, is something engineers tolerate as part of advancing technology (provided the engineer in question exercised reasonable care). What engineers do not tolerate is that an engineer, any engineer, should make the same mistake. Once a mistake has been identified, the state of the art advances and what was once tolerable becomes intolerable (a kind of incompetence). Because keeping good records is part of accountability, engineers have routines for recording what they do.[23] Indeed, anyone who watches engineers work will be surprised at how much of what they do is "documentation". Engineering is unusual among professions in recognizing an obligation to "acknowledge their errors."[24] (NSPE 2010, III.1.a.)

So, one response engineers have to the problem of many causes is to take responsibility for enough investigation to assign further responsibility (both for past errors and for future tasks). They seek to pierce the fog of technology rather than viewing it as a cover under which they can escape responsibility. The argument from many causes overlooks the possibility of that response—indeed,

of that sort of responsibility-as-liability (liability as a duty to investigate and report).

But once engineers have investigated, they will often—perhaps always—have at least three candidates for *the* cause (or the chief causes) of the same event: operator error (running the THERAC-25 contrary to instructions), organizational failure (inadequate training or supervision, inadequate quality assurance, or the like), and design flaw (faulty micro-switch or software bugs). There is a nice joke to illustrate how fundamental this trinity of candidate causes is:

> God says, "I told you not to eat from that tree. You have no one to blame but yourselves for your fall from grace."
>
> "But," Eve responds, "you should have explained better. You should have been more explicit about the consequences."
>
> "And," Adam adds, "just what was that tree doing in the Garden, anyway?
>
> Would it not have been safer to plant it where we could not go?"

God is trying to treat operator error as *the* cause of the disaster. He accepts no responsibility for what happened. He has a point, of course: Adam and Eve did what they knew they should not do. Had they done otherwise, there would have been no problem. They are, in a sense, plainly the cause of their own fall from grace. But Eve also has a point. Had God been more explicit about the consequences, she might have been better able to resist the serpent's enticements. God certainly could have done a better job of training (and supervising) his subordinates. The organization (God and his angels) failed to do all they reasonably could to prevent the disaster. But for that failure, Adam and Eve might still be in the Garden. And Adam also has a point. Had God designed the Garden differently, Eve could not have reached the tree. Why did God put the tree in the Garden? (And why, we might add, did God not fence the serpent out?) The Garden was a disaster waiting to happen—and God was responsible for that.

Generally, when something goes wrong with technology, there will be this trinity of candidate causes. That fact yields a version of the argument from many causes independent of ordinary complexity. This version relies instead on a fundamental problem in all talk of causes, that is, distinguishing mere causal factors, always many, from the cause or chief causes. What is a causal factor is (more or less) a simple matter of fact (or metaphysics).[25] What is *the* cause, or one of the *chief* causes, is not. As Joel Feinberg noted almost a half century ago, which causal factor we cite as cause depends on practical considerations, especially on the reason we are looking for a cause, what Feinberg called a "pragmatic criterion". Feinberg (1970) distinguished three such criteria (calling them): *explanation, engineering,* and *blame*. If our purpose is explanation, we look for the causal factor that helps us to understand why the event in question occurred. The serpent's presence in the Garden, and its cunning, may well be the best explanation of why Eve ate the apple. Without the serpent, the story would have been different. Adam and Eve might have lived in the Garden many happy years before they did anything wrong. They might be there still. The serpent was (in this respect) the cause of the fall from grace. Yet, if our purpose is "engineering" (in Feinberg's sense), then both Eve's organizational criticism and Adam's design criticism seem to pick out true causes. Both point to ways to prevent similar disasters. Eliminating the serpent would only eliminate "the immediate cause" (a particular tempter), not the "underlying" or "real" cause (as we might say), which is temping fruit, something better design or training might eliminate. If, however, our purpose is simply to assign blame, then, and only then, does God seem to have a point. Adam and Eve have no one but themselves to blame for their fall from grace. The serpent was not free in the way they were; the Garden did not eat the fruit, they did; and they knew better.

The practical criterion most appropriate to engineering as such is neither explanation (which seems to belong to history or science) nor blame (which seems more appropriate to law or ordinary moral life), but to what Feinberg actually called "engineering"[26]. The primary concern of engineers who have identified an engineering problem, whether or not of their own making, is—all else equal—to fix it if they can.[27] This is a corollary of the purpose of engineering (roughly, to improve the material conditions of humankind).

It is certainly the engineer's responsibility-as-domain-of-tasks. So, for example, though fools are to blame for much they suffer, a routine part of engineering is to make designs "fool proof", that is, to protect against foreseeable misuse (at least insofar as that is possible at reasonable cost).[28]

For engineers, then, the second version of the problem of many causes has a simple solution. Because the engineering criterion distinguishes cause from mere causal factor in a specific way, it forecloses the argument from many causes (in its second version). Indeed, engineers have constructed their profession to foreclose that argument. They make it their responsibility to solve certain problems, engineering problems. Among those problems is avoiding certain disasters (whether by design of the thing itself or by design of its organizational context). Solving such problems is something engineers are "liable to do" because they are engineers, not because they caused the problem. Here is a good example of (moral) responsibility-as-liability without responsibility-as-simple-cause.

This response to the argument from many causes may seem to dismiss the question of blame. It does not. Instead, it points out that engineers can, and often do, take responsibility where blame cannot be assigned. This response does, however, pose a question for those who find it unsatisfactory because it says nothing about blame. Why is blame so important here? Is it not more important to fix the problem, say, the malfunction in the THERAC-25, and to avoid similar failures in the future, than to assign blame for the original failure? (Perhaps engineers and the public will answer this question one way; lawyers and their clients, another.)

Replaceability

We turn now to the third argument against responsibility, the argument from replaceability. Like the first two arguments, this one does not depend on technology or any special feature of modern organizations. The argument is often used in some such form as this: "If I don't do this, someone else will, and since it makes no difwrong with me doing it." (Glover 1975; Scott-Taggart 1975; and Bayles 1979).Though that is the common form of the argument, both its strength and weaknesses will be clearer if we state the argument more formally:

1. To be the cause, or a chief cause, a causal factor must be necessary to the event caused.
2. If I do not bring X about, someone else will.
3. If X will occur whether I bring it about or not, I cannot be necessary for X.
4. If I am not necessary for X, I cannot be the cause of X or even a chief cause.
5. No one is responsible for what he did not cause.

Therefore: I am not responsible for X (whether I bring X about or not).

There are at least three flaws in this argument—at least as applied to engineers.[29]

The first flaw concerns premise 5. It claims that no one is responsible for what he did not cause. We have just seen that that is not true, that engineers can—and do—take responsibility for problems they did not cause (and are therefore responsible for them—in both accountability and liability senses).

The second flaw is factual. The claim that someone else will do X if I do not (premise 2) is true only if the agent in question can in fact be replaced by someone who will do X. Having long ago identified replaceability as a problem, engineers have adopted standards of practice, including a code of ethics, to prevent employers from being able to make the replaceability argument to engineers in a wide range of circumstances. Within that range, the engineer can simply say, "If you want to replace me with an unethical engineer, go ahead. There may be a few out there. But no ethical engineer will do what you are asking. Do you really want an unethical engineer working for you?" Engineers have in this way largely foreclosed the argument from replaceability. Normally, the argument will be unsound. Except in extraordinary circumstances, there will be no one out there both able and willing to violate the engineering standard in question. And, in those extraordinary circum-

person—that person being either an "unethical engineer" or no engineer at all.

The third flaw in the argument is logical. The argument assumes that there is only one criterion for cause when, as we have seen, there are at least three. Even if the engineer can be replaced, either with another engineer or with some other sort of technologist, the argument from replaceability carries weight only when we are considering what to fix. (If any operator would do X, the problem of X-ing is not solved by replacing this operator.) But to consider only what to fix is to rely on (a narrow version of) the *engineering* criterion of cause. We might, instead (or in addition), be concerned with finding the cause because we are seeking to avoid deserved blame, damages, or punishment, or having to give an account. We would then appeal to a different criterion of cause. We might, for example, say that X, an act such as approving a flawed design, is wrong even if the approval would be given by a few other engineers if the engineer in question did not give it. Whoever does X is at fault and that individual's fault is the cause of any harm that actually follows. The possibility of alternate causal chains does not change the actual causal chain or make the cause in question less faulty. (The only question is whose fault it will be, X's or someone else's, assuming it is someone's.) Premise 3 seems to equivocate between a first-person version (my doing y caused X) and an impersonal version (someone's doing y caused X). The argument can block responsibility for what "I" in fact do only by adopting an interpretation of premise 3 in which "I" disappears from the causal chain. The argument from replaceability thus begs the question—whether *I* should contribute to the causal chain in question.[30]

Institutional Constraint

The argument from institutional constraint, though another argument (more or less) independent of technology, is not independent of organization. The argument seeks to undercut the claim that engineers have a choice. Lynch and Kline state it nicely:

> Most engineers operate in an environment where their capacity to make decisions is constrained by the corporate or organizational culture in which they work. Engineers are rarely free to design technologies apart from cost and schedule pressures imposed by a corporate hierarchy, a government agency concerned with its image, or market pressures. (Lynch and Kline 2000, 210)

Lynch and Kline are certainly right that most engineers do (and always have) worked in large organizations. Except for those engineers high in the organization, working in such an organization means working within a framework that others have constructed, one that necessarily drastically limits what an individual engineer can do. Of course, the other side of working within a large organization is that one's choices, while limited by cost, market, organizational politics, and the like, can have much larger effects than the same sort of choice outside the organization. To design a small plane in one's own little company is unlikely to affect many people. To design just the riveting for the wing of the Boeing 737 is likely to have a much larger effect. For engineers, working for a large organization is a Faustian bargain—except that Mephistopheles, if present at all, is largely in the details.

Can an engineer have responsibility for what her employer does—or just for what she does within the narrow bounds where she is free? I see no reason why the engineer cannot take responsibility for what her employer does, however large the employer. She has voluntarily accepted employment with that employer, and can break off that association at any time just by giving notice. Certainly, engineers have long recognized that they should take responsibility for what their employers do—even when what the employer does is beyond the individual engineer's control. Among the provisions of one of the earliest codes of engineering ethics is this: "If after becoming associated with an enterprise he finds it to be of questionable character, he should sever his connection with it as soon as practicable" (American Institute of Electrical Engineers 1912, A.2.). Many contemporary codes contain a similar provision. At a minimum, an engineer should be willing to give an account of an employer's conduct and to accept some blame for what it does when what it does deserves blame (just as she

might accept credit if the employer does something creditable). An engineer who declines to accept that minimum responsibility but continues to benefit from the employment will seem much like that great but greatly flawed engineer in Tom Lehrer's song:

> "Once the rockets are up, who cares where they come down? That's not my department," says Wernher von Braun.

While engineers work *in* departments, they work *for* an organization. If the organization's name is on the rocket, they may have something to answer for when it comes down—even if their department has nothing to do with rockets.

That, however, is not the end of the argument from institutional constraint. The main point of the argument is to free engineers from responsibility for what *is* "their department", that is, for what they in fact help to bring about. The engineers are causal factors but (the argument runs) the organization so hems them in that they are not free enough to *be faulty* causes. The engineer who prepares a plan for a project is not responsible for it because she could have prepared no other. She had "no choice".

This version of the argument from institutional constraint has at least three flaws (beside the one already discussed). Each flaw is serious in itself. One is logical. The engineer always has at least one other choice, that is, not to prepare the plan. That other option may be suicidal, but—as a matter of logic—it always exists. Morally, there is an important difference between "no choice" and "no attractive choice".

The second flaw is conceptual. Engineers are hired to exercise engineering judgment on behalf of their employer. Where judgment is necessary, there must be at least two options, aside from preparing no plan. If there is only one option, a technician can prepare the plan. Even if engineers do not make the final decision, they must make the initial decisions on almost any project they undertake. That decision-making power, that freedom to choose, is what, as a conceptual matter, distinguishes an engineer from a mere technician. The argument from institutional constraint is not a reason why engineers cannot be responsible for

in the faulty-cause, accountability, or liability sense of responsibility. They have the necessary freedom to choose.

The third flaw is empirical. Those of us who have watched engineers at work have noticed that many of the decisions engineers make are *in fact* final. Superiors review them but, unless something is clearly wrong, the superior will not enquire deeply into the decision. Going along with engineering decisions is generally more efficient than getting a second opinion or overruling the engineer "on principle". As a matter of fact, *most* engineers in large organizations have considerable decision-making power. Whenever that is true, the argument from institutional constraint simply does not apply.[81]

Individual Helplessness

But, it may be said, the problem is not so much freedom as the need to cooperate with others to accomplish the task in question. One engineer cannot, for example, do anything to help her company adopt sustainable practices that would raise the price of the product in question. The market blocks the way. The individual engineer is helpless. Hence, even though the engineer's code may, as the American Society of Mechanical Engineers or American Society of Civil Engineers code of ethics does, require engineers to "strive to comply with the principles of sustainable development in the performance of their professional duties", an individual engineer can in fact do nothing. Lacking any power to follow the code in this respect, the engineer can have no responsibility to do so. See for example, Miller (2005).

So long as we think of engineers as individuals without the power to coordinate their conduct, this argument from helplessness is sound—at least for a large number of engineering decisions. The flaw in the argument is that engineers, as members of a profession, are never mere individuals. They are always, in addition, members of the profession. The profession may adopt standards of practice which, applying across all employers in a market, would eliminate market pressures that might otherwise block adoption of the practice. Where the market still resists adoption of the standards, the engineers may, through

companies, other private entities, or the government to enforce the standards. (Which entity should be appealed to will, of course, depend on who would benefit from the standard's adoption.) American engineering societies have been developing technical standards of this sort for more than a century. There is a more recent history of similar standard-setting by international bodies, such as the IEEE or International Organization for Standardization (ISO).[32]

Individual Ignorance

But, it may be said, the problem is not so much freedom or power as knowledge. The complexities of any large project make it hard for most engineers to know the whole. On many large projects, engineers literally do not know what they are doing. And, because they do not, they cannot be responsible in either the liability or accountability sense for what happens. That is the sixth argument against responsibility, the argument from individual ignorance.[33] It has at least two flaws.

The first flaw should be obvious from what we have already said. Engineers generally try to understand how what they are working on fits into the world. That is part of good engineering. An engineer who literally does not know what she is doing has a lot to account for. There may, it is true, be some environments where engineers are kept ignorant of the big picture. Some military work is said to be organized to keep most participants in the dark about the ultimate purpose. And some projects, especially those with large software components, are unusually hard to understand. Like the THERAC-25, they end up behaving in ways that no one foresaw—some, perhaps, in ways no one could foresee given the state of the art at the time. But even such unusual environments leave engineers with a responsibility to try to find out what they do not know but need to know to do a good job; it does not excuse them from trying to do a good job. The argument treats the elimination of some responsibility-as-liability as the elimination of all liability.

Ignorance of outcome is, however, rare—and, in any case, only frees engineers from blame. That is the second flaw in the argument. Whatever force will still owe an accounting for what happened. And, if anything went wrong, they will also have a responsibility (a liability) to figure out how to prevent it from happening again. The argument from ignorance only defends against one sort of responsibility-as-liability, the sort engineers as such are least interested in (blame).[34]

Neutrality

That brings me to the last of my seven arguments, the argument from technology's neutrality. While it is true that engineers produce technology, they cannot, it is said, control what is done with it. Any piece of technology, even something as innocuous as a scissors, can as easily be put to a bad use (stabbing an innocent person) as a good one. Engineers should not be held responsible for what is done with their work.[35]

This argument does not, it should be noted, deprive engineers of all responsibility, for example, responsibility for flaws in design or manufacture over which they actually have control or, indeed, responsibility for uses for which the technology was designed. All it can deprive engineers of is responsibility for *misuses* of their work *beyond their control*. And, depending on the kind of liability or accountability in question, it may not be able to do even that. Consider the THERAC-25 disaster again. While engineers are certainly not to blame for what they could not control, their response to the disaster was not to crow about how they were not to blame. Instead, they viewed themselves as having a professional responsibility to prevent similar disasters now that they knew what in particular had gone wrong with the machine. For engineers, what they create should not be neutral. It should be designed, built, maintained, and even disposed of in a way at least consistent with the public health, safety, and welfare. Technology as such may, or may not, embody certain values. That is a deep question for the philosophy of technology.[36] *Engineered* technology, in contrast, necessarily embodies (or at least should embody) certain values. An engineer who produces a product that, however clever, fails to benefit humankind has failed—as an engineer. Engineers have been making that point about

separated from military engineering early in the nineteenth century (Davis 1998, pp. 12–15). The argument from neutral technology simply overlooks the distinction between technology as such and engineered technology. For that reason, as much as because of government regulation, today's scissors generally lack the sharp tips they once all had. Part of good engineering is designing to prevent foreseeable misuse. Another part of good engineering is seeking to add to the misuses that can be foreseen, for example, by tracking products in use to discover how they have been misused.

THE RATIONALITY OF TAKING RESPONSIBILITY

I have canvassed seven arguments seeking to relieve engineers of moral responsibility for what they do (or fail to do). These seven are, I believe, all the arguments now in the literature. I have dealt with them in different ways. Some I disposed of by detailed analysis of the key concepts (such as "cause") or by pointing to logical mistakes. Some I could dismiss because they relied on a claim for which there was substantial adverse evidence. But, for most, one step was to point out how engineers have taken on the responsibility in question (liability or accountability). The responsibility in question arises from their own voluntary acts (whether as individuals or as the engineering profession). Though people are generally thought to shy away from responsibility-as-liability or as-accountability, engineers do not seem to. Instead, they seem to claim certain responsibilities most of us, even technically trained managers and other technologists, try to avoid.[37] Some might well wonder whether I have described a mass pathology rather than explained why engineers can have such moral responsibilities. Why take on responsibilities others do not want?

That question, one concerned with rational justification, is one I must now answer. It is, I think, a question easily answered in a way preserving the moral responsibilities in question: engineers, as a group, gain more by taking on the responsibilities in question than they lose. Consider a joke engineers like to tell:

> A priest, a lawyer, and an engineer are about to be guillotined in one of the great squares of Paris during the Reign of Terror. The condemned must lie on a bench with his head set in braces under the huge angled blade. The priest is first. Asked whether he would like to lie face up or face down, he chooses face up—adding, "I want to look where I hope to go." He says a prayer and lies down. The blade drops—but just as it reaches his throat, it stops. The executioner declares "divine intervention", unties the priest, and lets him go. The lawyer is next and, given the same choice, chooses to "follow precedent—and hope for the best." Again the blade stops just short. The lawyer is also let go, but this time the executioner checks the machine before proceeding. He finds nothing amiss. The engineer is brought forward. He too chooses to lie face up. Just as the blade is about to drop, he shouts "*Stop!* I think I see the problem!"

This is a good joke and, like most good jokes, has several uses. I first heard it from an engineer making the point that not every engineering problem should be solved. I have also heard it used to make the point that engineers see problems others do not. Here, however, I tell it to illustrate how and why engineers might take on responsibilities not otherwise theirs. The engineer under the blade was not (as far as we know) even a causal factor in the production, installation, or maintenance of the guillotine, or under contract or otherwise obliged to take on its repair or even to investigate its malfunction. He might have kept silent without blame—and thereby saved his neck. Fixing the guillotine definitely was not "his department".

Yet, whatever we think of this engineer's prudence, we recognize in him someone we can trust with the world's machinery, someone quite unlike the hobo in the henhouse. Avoiding responsibility has many advantages, but winning the trust of others is not one of them. What engineers gain by taking on responsibilities others avoid are several tasks or offices which, as a matter of fact, have become a relatively lucrative occupation. Creating a lucrative occupation is rational enough to justify accepting the moral responsibility (both accountability and liability) making it possible (provided the tasks accepted are also morally permissible—as they seem to be). The very barriers to moral responsibility that technology and organization throw up are

opportunities for a profession willing to make overcoming them their responsibility.[38]

NOTES

1. See for example, Nissenbaum (1996). Though her focus is, as her title suggests, computer scientists, her argument fits engineers as well—and some of her examples include engineers.
2. Though responsibility-as-good-cause seems the mirror image of responsibility-as-faulty-cause, it is overlooked in most discussions of responsibility—perhaps because there is no obvious credit version of responsibility-as-competency responsibility-as-liability, or responsibility-as-accountability. The term for this missing sort of responsibility might be "responsibility-as-reward" (or perhaps "creditability").
3. "Responsibility-as-power" often appears in discussions of professional ethics. See for example, Alpern (1983), for an early example applied to engineering.
4. Hart (1968, 210–223) seems to have been the first to distinguish four of these senses: responsibility-as-role (which I have divided into "office" and "task"), responsibility-as-simple-cause, responsibility-as-liability, and responsibility-as-competency ("capacity"). Ladd (1982, esp. 64–65) seems to have been the first to distinguish accountability (which, unfortunately, he treats as a merely organizational notion—as if a perfect stranger cannot "call me to account" for something untoward I did as a mere individual). Kutlik (1999, esp. 174–175) offers six senses—some significantly different from mine. For example, he distinguishes responsibility-as-simple-cause (using a hurricane example) from responsibility-as-function ("the heart is responsible for pumping blood"). Though a useful distinction for some purposes, it would only complicate discussion here. Nevertheless, that distinction is a reminder that my list is neither exhaustive nor canonical.
5. Those who distinguish between "backward-looking liability" and "forward-looking liability" seem to me to have confused responsibility-as-liability with responsibility-as-simple-cause or responsibility-as-faulty-cause. The reason for that confusion seems to be that blame is (generally) backward-looking. We generally blame people for what they have done, not for what they will do, and liability (it is thought) must work in the same way. The mistake is to assume that other blame-related concepts must be just as backward-looking someone for doing such-and-such is to say of her that she *did* something she should not have (or failed to do something she should have), that is, to refer to the past, to say of someone that they *deserve* blame (or that they are blamable) is not only another way to blame them; saying that also sets out a set of tasks (blame them in various ways until the reason for blaming is gone), that is, says something about the future (as blaming does not). Deserving blame differs from responsibility-as-liability only in not assigning the tasks to anyone, though it gives each of us a reason and a right to do the tasks.
6. There should be a corresponding problem about engineers and good causation, but no one seems to want to deny engineers credit for their good works even though the same arguments that undermine liability and accountability would seem to undermine credit. Where there can be no blame, there can be no praise (or so we generally think).
7. Indeed, holding *persons* responsible without fault is not the limit of what the law can do. The law can also hold non-humans responsible. For example, in 1595, in Leiden, Provetie, an ordinary dog, was tried for the murder of a child whose death it caused. The dog nipped the child's hand while taking meat from it, the wound became infected, and the child died as a result. All this having been proved, Provetie was found guilty and publicly hanged. There was no claim that Provetie intended to harm the child, that it foresaw the harm, or that it even should have foreseen the harm. It was enough to hold the dog responsible for the death that it did what it should not have, nipped the child's hand, however innocently, and the child died as a result (Arthur and Shaw 2006, 245–246).
8. So, for example, morality is silent concerning the legal punishment of dogs. Since reason no longer seems to recommend such punishment, we may wonder why it ever did (if it did). One answer at least is obvious: Going through the forms of criminal trial may have been the simplest way to deal with an unusual case. There is some evidence that this obvious answer may also be the right one. The official summary of the case includes language probably standard at the time but wholly inappropriate to this particular case, for example, "all of which appears from the prisoner's own confession, made by him without torture or being put in irons" Arthur and Shaw, 246. How could the

9. Ladd (1982), 67. Note that Ladd explicitly connects moral responsibility with power (my fifth sense of responsibility), for example: "one of the principal factors that creates moral responsibilities for one person rather than another is a difference of *power*, which usually consists of superior knowledge and ability to affect outcomes."
10. Those who are not competent (in the appropriate respect) cannot (it is assumed) be at fault (with respect to that). Neither can those who could not have foreseen the harm they caused be at fault in that respect. Competence and knowledge (of the appropriate sort) are (it is assumed) at least preconditions of fault.
11. For those who think that I should have included "control" among the list of necessary conditions, I recommend Sher (2006). As will be plain, nothing important turns on this omission.
12. Her inaction is, in this respect, quite different from the inaction of a lifeguard assigned to watch a stretch of beach. His inaction is a causal factor in any drowning death he might have prevented precisely because another lifeguard would have been there had he not been. The beach would not (we may suppose) be open for swimming if a lifeguard were not present.
13. Lawyers like to interpret "cause" as necessary condition ("but for"). The problems with so interpreting "cause" are well-known (for example, an inability to make sense of redundant causes or to limit the number of causes). We need not concern ourselves with those problems here. The objection made here stands even if those problems can be solved. For recent work on omissions as causes (more or less consistent with what I say here), see: McGrath (2005), Dowe (2004), Boniolo and De Anna (2006), Pundik (2007), and Baumgartner (2008).
14. Curtis Forbes tells me that this comment may reveal me to be an urbanite. He has noticed that those who grew up in the countryside, as he did, are much less likely to feel this responsibility. He may be right, but even if this sense of responsibility (and the moral judgment of "worse" it supports) is "cultural-centric", it relies on a conception of responsibility independent of choice (and does not appear to be incoherent because it does), which is all I need to make my point.
15. Now, it might be objected that this example is really about a duty to aid because one is in a position to do so, rather than a duty to aid because one caused the problem. I do not think there is a think there are good reasons to dismiss it. The most important is that the duty to aid usually involves serious danger to another—not, as in this case, merely cleaning up a mess. The only people who are likely to see much force to this duty to aid, apart from the causal connection, are act utilitarians, supposing I have no better way to use my time than looking after the dead cat.
16. I am not the only one to notice this sort of responsibility. See for example, Kutz (2002), p. 558: "Sometimes an agent's mere causal linkage with harm may warrant a response....agents can reproach themselves for faultless conduct that causes a harm, even when their victims, and onlookers, do not reproach them." The first half of Kutz (2002) is worth reading for the way his "Strawsonian" interpretation of responsibility tends to reach the same conclusions I do in this section.
17. I am not the only one to notice this sort of responsibility. See for example, Kutz (2002), p. 558: "Sometimes an agent's mere causal linkage with harm may warrant a response ... agents can reproach themselves for faultless conduct that causes harm, even when their victims, and onlookers, do not reproach them." The first half of Kutz (2002) is worth reading for the way his "Strawsonian" interpretation of responsibility tends to reach the same conclusions I do in this section.
18. There is, I think, a close analogy between this view of responsibility as primarily imposed and what is called "the external" perspective on law. The external perspective on law misses essential features of law that the internal perspective includes, especially its authority. See especially, Hart (1961), pp. 101–102. From the outside, responsibility may seem to be primarily about having to answer to others for what one does (what Kant would call "heteronomy"). From the inside, responsibility may seem to be primarily about being trustworthy, someone capable of answering for what he does, a certain (virtuous) disposition (what is recognizably Kant's "autonomy").
19. Nissenbaum (1996, 29).
20. I am inclined to think that discussions of responsibility suffer from supposing, however implicitly, too close an analogy between moral and legal responsibility. While legal responsibility is primarily about holding others responsible, moral responsibility is at least as much about holding ourselves responsible as holding others responsible. The structure of moral responsibility may (despite many overlaps with legal responsibility) have far

even blame than legal responsibility has. For more on this, see Cane (2002).

21. Nissenbaum, 30, has therefore confused the problem of many hands with the problem of many causes. She is not the only one. See for example, Bovens (1998, esp. 45–50), a section titled "Accountability: the problem of many hands" and beginning with the Thompson quotation used above, but consisting of several interesting examples all of which illustrate the problem of many causes rather than many hands.
22. The failure to take account of such differences between professions is one reason why Bovens' treatment of professional responsibility (Bovens 1998, 161–163) is—like that of most social scientists and the philosophers who follow their lead—decidedly misleading.
23. The obligation to keep good records, though helpful to accountability (should giving an account become necessary), is not itself accountability but a liability (an obligation to do something other than give an account).
24. Some other engineering codes have similar provisions. See for example, IEEE (2006): "7. to seek, accept, and offer honest criticism of technical work, to acknowledge and correct errors, and to credit properly the contributions of others." (I cite codes of engineering ethics in this paper to demonstrate that I am not making up what I claim to be facts about engineering. The codes constitute evidence, though not decisive evidence, of what responsibilities engineers have taken on.)
25. Well, relatively simple. There are conceptual problems about how to treat redundant causes, alternate causes, and the like (as noted already)—and these affect (and are affected by) how we conceptualize cause itself. See especially, Hart and Honoré (1973).
26. Of course, engineers do blame one another sometimes, but there is nothing especially interesting about that. Their blaming will have much the same structure as ordinary blaming—though it will take into account the special responsibilities engineers have taken on.
27. Feinberg (1970), 205, eventually calls the engineering criterion "the handle criterion", having in view the expression "to get a handle on it".
28. Whether this is a good idea is a separate question. Compare Bucciarelli (2002).
29. There is a fourth flaw which raises deeper questions about causation. Premise 1 also seems to be false because redundant causes can be causes. Suppose two hunters shoot the same bystander at the same moment (because he has come between them and a duck they are both aiming at), that both their bullets enter his body at the same time, and that either bullet would be sufficient to cause death. Neither hunter is necessary for the bystander to die. The death is "over-determined". Yet, it would seem odd to let both off because neither was the cause, or even a cause, of the bystander's death. I ignore this flaw (and several related ones) here because I promised charity on the subject of necessary conditions for responsibility.
30. For a more detailed discussion of this argument, see the exchange: Davis (1986a, b) and Bayles (1986).
31. Swierstra and Jelsma (2006) draw this conclusion from their study of European engineers. I came to the same conclusion from a study in the United States. See Davis (1997).
32. So, for example, a claim like the following seems to me clearly mistaken: "As an engineer, I think we take ourselves too seriously if we think we are ever going to influence what society decides to do by looking at the problem of ethics within the context of technology itself" Wiseman (1980), p. 166. This argument has an extension to society as a whole not relevant here. See for example, Winner (1995).
33. See for example, Florman (1978), 323: "An engineer designing a rapid transit cannot become expert in acoustics, urban planning, and the habits of woodland birds, and at the same time be an expert in the design of monorails. Nor can he do his best work if he is excessively apprehensive about the consequences of his every move."
34. Of course, engineers are interested in avoiding blame, just as they are interested in avoiding legal liability. All I claim is that, as engineers, that is (and should be) what they are *least* concerned about—just as lawyers, as such, are (or at least should be) most concerned with legal liability. As individuals, avoiding blame may be their chief concern.
35. See for example, Schnädelbach (1980), 28: "'We dealt only with the technical problems and had no influence on the determination of goals'—this is a type of excuse frequently advanced in order to separate technical from ethical responsibility after political and moral catastrophes." While Schnädelbach does a good job of stating the argument from neutrality, he in fact rejects it—even for applied

science. I am still looking for a contemporary thinker who endorses it.

36. For a good survey of this debate, see Sundström (1998). By "technology", I simply mean all those systems of people and things that constitute the world humans have made, everything from bridges, computer programs, and dictionaries to hybrid corn, hot dogs, and Labrador retrievers.
37. Think, for example, of programmers who claim to be "artists rather than engineers"—or scientists who rely on the neutrality argument discussed in the last section.
38. Compare Ladd (1982), 66: "responsibilities are not incurred or acquired like obligations. Rather one finds oneself responsible for something or other as a result of being in a certain position, e.g. of power." The responsibility the engineer took here is (Ladd notwithstanding) incurred like an obligation, not by having the power to help but by a specific voluntary act, claiming to "see the problem". Had the engineer remained silent, he would have had no responsibility to fix the guillotine.

REFERENCES

Alpern, K. (1983). Moral responsibility for engineers. *Business and Professional Ethics Journal*, 2, 39–48.

American Institute of Electrical Engineers. (1912). *Code of ethics* (March 8). http://ethics.iit.edu/publication/CODE-Exxon%20Module.pdf. Accessed July 22, 2010.

Arthur, J., & Shaw, W. H. (Eds.). (2006). *Readings in the philosophy of law* (4th ed.). Upper Saddle River, NJ: Pearson Prentice Hall.

Baumgartner, M. (2008). Regularity theories reassessed. *Philosophia*, 36, 327–354.

Bayles, M. D. (2008). Regularity theories reassessed. *Philosophia*, 36, 327–354.

Bayles, M. D. (1979). A problem of clean hands. *Social Theory and Practice*, 5, 165–181.

Bayles, M. D. (1986). Reply to Davis. In Davis & Elliston (1986) (pp. 458–460).

Boniolo, G., & De Anna, G. (2006). The four faces of omission: Ontology, terminology, epistemology, and ethics. *Philosophical Explorations*, 19, 277–293.

Bovens, M. (1998). *The quest for responsibility*. Cambridge: Cambridge University Press.

Bucciarelli, L. L. (2002). Is idiot proof safe enough? *Applied Philosophy*, 2, 49–57.

Cane, P. (2002). *Responsibility in law and morality*. Oxford: Hart Publishing.

Davis, M. (1986a). The right to refuse a case. In Davis & Elliston (1986) (pp. 441–457).

Davis, M. (1986b). Rejoinder to Bayles. In Davis & Elliston (1986) (pp. 451–464).

Davis, M. (1997). Better communications between engineers and managers: Some ways to prevent ethically hard choices. *Science and Engineering Ethics*, 3, 171–213.

Davis, M. (1998). *Thinking like an engineer: Essays in the ethics of a profession*. New York: Oxford University Press.

Davis, M., & Elliston, F. (Eds.). (1986). *Ethics and the legal profession*. Buffalo, NY: Prometheus Books.

Dowe, P. (2004). Causes are physically connected to their effects; Why preventers and omissions are not causes. In C. Hitchcock (Ed.), *Contemporary debates in philosophy of science* (pp. 189–196). Oxford: Wiley-Blackwell.

Feinberg, J. (1970). *Doing and deserving*. Princeton, NJ: Princeton University Press.

Florman, S. C. (1978). Moral blueprints. *Harper's*, 257, 311–323.

Glover, J. (1975). It makes no difference whether or not i do it: Part 1. *Aristotelian Society Supplementary Volume*, 49, 171–190.

Hart, H. L. A. (1961). *The concept of law*. Oxford: Oxford University Press.

Hart, H. L. A. (1968). *Punishment and responsibility*. New York: Oxford University Press.

Hart, H. L. A., & Honoré, A. M. (1973). *Causation in the law*. Oxford: Oxford University Press.

IEEE. (2006). *Code of ethics*. http://www.ieee.org/membership_services/membership/ethics_code.html. Accessed July 22, 2010.

Kranzberg, M. (Ed.). (1980). *Ethics in an age of pervasive technology*. Boulder: Westview Press.

Kuflik, A. (1999). Computers in control: Rational transfer of authority or irresponsible abdication of autonomy? *Ethics and Information Technology*, 1, 173–184.

Kutz, C. (2002). Responsibility. In J. Coleman & S. Shapiro (Eds.), *The Oxford handbook of jurisprudence and philosophy of law*. Oxford: Oxford University Press.

Ladd, J. (1982). Philosophical remarks on professional responsibility in organizations. *The International Journal of Applied Philosophy*, 1, 58–70.

Lynch, T., & Kline, R. (2000). Engineering practice and engineering ethics. *Science, Technology, & Human Values*, 25, 195–225.

McGrath, S. (2005). Causation by omission: A

Miller, D. (2005). Distributing responsibilities. *The Journal of Political Philosophy, 9*, 453–471.

Nissenbaum, H. (1996). Accountability in a computerized society. *Science and Engineering Ethics, 2*, 25–42.

NSPE. (2010). *Code of ethics for engineers.* National Society of Professional Engineers, http://www.nspe.org/Ethics/CodeofEthics/index.html. Accessed July 22.

Pundik, A. (2007). Can one deny both causation by omission and causal pluralism? The case of legal causation. In Russo & J. Williamson (2007) (pp. 379–412).

Russo, F., & Williamson, J. (Eds.). (2007). *Causality and probability in the sciences.* London, UK: College Publications.

Schnädelbach, H. (1980). Is technology ethically neutral? In Kranzberg (1980).

Scott-Taggart, M. J. (1975). It makes no difference whether or not I do it: Part 2. *Aristotelian Society Supplementary Volume, 49*, 191–209.

Sher, G. (2006). Out of control. *Ethics, 116*, 285–301.

Sundström, P. (1998). Interpreting the notion that technology is value-neutral. *Medicine, Health Care and Philosophy, 1*, 41–45.

Swierstra, T., & Jelsma, J. (2006). Practice responsibility without moralism in technoscientific design. *Science, Technology, and Human Values, 31*, 309–332.

Thompson, D. E. (1987). *Political ethics and public office.* Cambridge, MA: Harvard University Press.

Winner, L. (1995). The enduring dilemmas of autonomous technique. *Bulletin of Science, Technology & Society, 15*, 67–72.

Wiseman, G. (1980). The powerlessness of engineers. In Kranzberg (1980).

1.3.3 TECHNOLOGICAL SUBVERSION

DAVID STRONG

In this reading taken from the author's book *Crazy Mountains: Learning from Wilderness to Weigh Technology,* philosopher David Strong asks us to consider the basic question, "Why do we value technology?"

In his analysis, we value technological devices because they disburden us of toilsome labor, discomfort, and bother. Technological innovations beckon us with the promise of freedom and happiness. We want things to be easy so that we are freer to pursue our ends without having to worry too much about the means we use to do so, and we believe that being "free" in this sense—that is, disburdened—will make us happy.

Using as his examples the replacement of the hearth by the central heating system and television viewing as a primary form of entertainment, Strong argues that this quest for freedom and happiness through technology provides ease and comfort, often alters the nature of human experience in ways that disconnect us from nature and other people, and ironically produces not happiness but alienation. Although some readers may find this philosophical meditation somewhat demanding, those who study it carefully will find that it rewards their effort.

Source: Reprinted by permission from *Crazy Mountains: Learning from Wilderness to Weigh Technology* by David Strong, the State University of New York Press © 1995, State University of New York. All rights reserved.

FOCUS QUESTIONS

1. What does the author mean by the term *technological availability*? What specific values does this concept embody?
2. What is the vision of the "good life" that our modern technological society offers us? Why does the author question the goodness of this way of life?
3. What are the main differences between what Strong calls "things" and what he calls "devices"? What values do we sacrifice when we choose the device paradigm over engagement with things?

KEYWORDS

alienation, commodification, environmental ethics, heating systems, television, wilderness

THE UNDERLYING ETHIC OF TECHNOLOGY

Some have argued that we live in an invisible iron cage. Indeed, technological forces are shaping people's lives in ways that they have little or no control over, especially if the basic framework of technology goes unchallenged; but, as Charles Taylor points out, the conquest of nature had a benevolent point to it. It was to serve humanity. So, he finds that, along with other forces, there are moral forces of work here shaping our lives. We live neither in an iron cage nor in an arena of unconstrained choice; we inhabit a possibility space where some moral choices are being made. There is a kind of ethical appeal to not letting our resources go to waste. So what ethical forces might be called upon to reform technology in a deep way? How should we understand the basic choice we face? For developing what I call the vision and underlying ethic of technology, I will draw heavily upon Albert Borgmann's theory of technology, the best account of the character of the technological culture we have so far. Then we will use this vision of technology to show that the concerns of environmental ethics and people's better concerns for nature generally will be subverted by technology unless we as a culture come to grips with the irony of this vision and begin to make a fundamentally different choice, that is, choose things over consumption.

MAKING THE APPEAL OF TECHNOLOGY INTELLIGIBLE: THE PROMISE OF TECHNOLOGY

Neither Heidegger nor Thoreau makes clear what it is about technology that is attractive to people. Claiming that we delight in the exercise of power seems correct enough when we think of the enormous amount of power we wield with technology, yet this view does not address our more intelligible motives and, therefore, does not really address many of the proponents of technology without trivializing their concerns. In one way or another, most of us, if not all, see technology as good. What is at the heart of our petty homocentrism? What good is technology?

Typically people articulate what good technology is when they say that something is better or improved and demonstrate that "that's progress." Advertisements are continually pointing out what is better about the product advertised, even if the chief "advantage" is two for the price of one. Although they may well dupe us, these advertisements normally appeal to standards that at least on a deep and general level are already in place and widely shared in consumer culture. We hear everywhere around us, not just in advertisements, what better is. "It means less work." "It's more comfortable." "It's convenient." "It's healthier." "It's faster and more productive." "It's less of a hassle." "Sleeker looks better." "It's lighter." "It doesn't get in your way." "You don't

have to wait on anyone else." "It's exciting." When we see the very latest devices, often our expressions are on the order of "Wow!" or "That's great!" or "Look at that, would you?" So, at deeper levels, there seems to be a good deal of like-mindedness about what constitutes better in our culture.

Another approach is to consider what people think of as clear examples of progress. Television today is far different from what it was in the past. In the early 50s, one was lucky to own a television. Reception was poor, the picture rough and in black and white; the screen was small, the set large; the number of programs was very limited. In addition to other obvious improvements, now the sets come on instantly, are controlled from the couch, can be found in all sizes and nearly everywhere, and have access to a vast number and variety of programs, especially with video cassette recorders. Even if they are not willing to pay the price for all of them, most count these changes as improvements, and rarely do we find people watching a black and white set any longer. What are the standards which make these changes count as improvements?

Television as a clear example of technology will play a key role in our understanding of the nature of the fundamental choice we face, but Borgmann uses another paradigmatic example of technology, the central heating system, to disclose most of these standards of technology. We can easily trace the development of central heating systems back to the wood-burning stove or the hearth. The chief advantages of the heating system over these latter two are various. Central heating *is easier*. We do not have to gather, stack, chop, or carry the wood. An automatic thermostat means that we do not have to trouble ourselves in the morning or evening with setting a thermostat. Central heating is more *instantaneous*. We do not have to wait for the house to warm up. It's *ubiquitous*. Warmth is provided to each corner of the room, to every room, and everywhere equally well. Finally, a central heating system is *safer* than a hearth. My grandmother was born in a newly built chicken coop because three weeks earlier her family's house burnt down from a chimney fire. So the standards by which people judge central heating to be better than a wood-burning stove are ease, instantaneity, ubiquity, and safety or some combination of these, for example, convenience. These four "technological standards" can be collected under the more general notion of technological availability. To be more *available* is to be an improvement, then, in terms of one or more of the four above standards.

Why does it seem to people that this availability is good? From one perspective, this availability relieves people of burdens: less effort, less time, and less learning skills are required. Available anytime and anyplace, they are disburdened of the constraints of time and place. They are disburdened of having to take risks. Historically, modern technology was envisioned as enabling people not just to subjugate nature, but to do so for the purpose of freeing humanity from misery and toil. To be relieved of these burdens then fulfills this vision of technology. *To the degree people personally share this vision*, they will also see its concrete manifestations, such as central heating, as unquestionably good. Compared to older versions, the latest portable computers exemplify this relief from burdens and are attractive to many for this reason.

By overcoming nature, technology would, as some in the seventeenth century foresaw, not only relieve humans of burdens, but it would make available to them—easier, safer, quicker, and more ubiquitous—all the goods of the Earth. So, technological availability negatively disburdens people of misery and toil, and positively enriches their lives, makes them happy, it seems. So seen, technology has an attractive glow about it.

> Technology promises to bring the forces of nature and culture under control, to liberate us from misery and toil, and to enrich our lives ... [More accurately], implied in the technological mode of taking up with the world there is a promise that this approach to reality will, by way of the domination of nature, fuel liberation and enrichment.

Borgmann calls this "the promise of technology."

Clearly those below the middle-class of advanced industrialized countries and those outside those countries do not derive the benefits of

technology, although many do feel the pull of its promise. The claims of social justice will not likely be met until the more privileged ones, the middle and higher classes of these industrialized countries, come to terms with the questionable character of technology's promise. So, the critique of technology I am developing here does not apply to those in poverty. It applies only to those who have too much.

For these latter, technology has made good on its promise in important ways. My grandmother's father died from what she believes was pneumonia when she was eleven, leaving her and her younger sister to perform heroic feats to save the cattle from starvation in the drought times of an extended winter that followed. Often hitching the team up before dawn and returning hours after dark, especially in winter, her family took an entire day to get to and from town sixteen miles away. For the privileged, then, many past hardships have now been conquered. Although we may have legitimate concerns about whether there is too much medical technology, no one could reasonably refuse every advance of modern medical technology. The weather will never be brought under control, but, via comfortable structures, nature's heat and cold, rain and snow are controlled as well as darkness and drought. Toilsome labor is largely eliminated within the culture of technology.

A reasonable person may reject motorcycles in favor of horses to do ranch work, but that person still rides to town on paved roads in a car, has parts shipped by air, reads a newspaper and books, transacts business over the phone, and owns at least a radio. No thoughtful person will want to turn her back on technology entirely. Thus, technology, by conquering nature, has relieved humans of severe burdens. Today we are still working to overcome those, such as cancer and AIDS, that remain. So, if technology does not saddle us in the long run with more than it has relieved us from, it will have made good on this aspect of what at first seemed and still does seem promising about it. It could turn out that ozone depletion, global warming, ecosystem destruction, the population explosion, polluted land, air, streams and oceans, and human and mechanical errors will impose burdens far greater than those we were relieved from in the first place. To meet these problems certainly calls for a reform of present practices. We read or hear of these calls for reform nearly every day. More common critiques of technology, such as David Ehrenfeld's *The Arrogance of Humanism,* attempt to show that technology will fail its own standards, bringing disaster upon us.

Much as reform in these areas is needed and much as these pessimistic critiques deserve thoughtful consideration, the present work will turn to a uniquely different task. It grants and, in fact, seeks to have the reader appreciate the genuine success of modern technology. Technology has relieved, and technology will, I assume for the purposes at hand, continue to relieve, humans of many hardships of the human condition.

So what is wrong with technology for those within the realm of its benefits? Underlying these standards of availability is really a vision of a good life that is free and prosperous. What is at the bottom of concern with technological availability is an aspiration for freedom and happiness. Most people, at least in the Western tradition, are concerned with liberty and prosperity. For Aristotle only the Greek free man was able to have sufficient time and sufficient wherewithal to develop the moral and intellectual virtues he thought to be required for happiness or eudaemonia. The Hebrew people's understanding of the covenant centered on an idea of prosperity. Jesus preached of a free and abundant life. The Enlightenment, as we see its results in "the pursuit of happiness" in our Declaration of Independence, is fully within this tradition. But to find agreement at this high level of abstraction is not to see that the crucial differences lie with the particular versions of freedom and prosperity. For Socrates living well had to do with human excellence and living a just life, not with materialism. The blessed life and the abundant life of the Hebrews and of Jesus was not commodity happiness. So, too, we must look carefully at the particular idea of freedom and prosperity governing people's attraction to technology, for only at this level of particularity will its misleading and harmful features begin to show. In other words, one can criticize the

trivialized forms of freedom and prosperity on which the technological society is centered without, at the same time, criticizing freedom and prosperity more generally as a vision of the good life. Quite the contrary, we can call technology into question even more sharply by showing that technology fails to provide the free, prosperous, and good life we want in our waking moments.

Technological society offers a flattened vision of freedom and prosperity. The more disburdened, the better off I am according to this vision. So, the technological idea of freedom is really one of disburdenment. What about prosperity? Cellular phones are currently a status symbol. These devices which disburden people of the constraints of place are taken to be a sign of affluence because, generally, only the more prosperous have them. So, in part, to be prosperous is to have the latest, most refined device. A sign of affluence, too, is to be able to go to an undiscovered exotic place, have the most channels and compact discs, own specially designed clothing, own what no one else has yet. Thus, in part, to be prosperous is to own the most varied, the widest assortment of commodities. Finally, when people buy a product on sale they get both the commodity they purchased and still have money left over. Why is that attractive? Because they can buy something else with the money saved. They are better off that way, they think, because they get more items for the money. Thus people pursue prosperity through the standards of owning the *most numerous, widest variety, and the very latest (most refined)* commodities. The powers that be in the technological society own and control most of these items. Such is the picture of the good life envied by those keeping up with the Joneses. Our culture's vision of the good life is the goods life.

Does this vision really deliver a good life? If we say no merely because it differs from the blessed life according to Abraham, Moses, and the prophets, or from the Greeks' eudaemonia, our analysis would be dogmatic and presumptuous. Technology must be thought through; it will not be met by simply reacting against it. So, if we answer "no," as I will, then we must be able to provide good reasons.

THE TECHNOLOGICAL MEANS TO FREEDOM AND HAPPINESS: THE DEVICE

The ironic consequences of this vision of freedom and prosperity can be drawn out through a careful analysis of the peculiar way technology transforms, or more specifically, dominates nature and culture. Technology does not dominate these in the traditional manner of lording it over them; rather, as Albert Borgmann shows, technology follows a pattern, unique to the modern era, in the way it gets everything under control. We can expose this pattern by examining instances of it.

The central heating system dominates warmth; it brings warmth under control in ways that the wood-burning stoves do not. To show its unique form of domination, Borgmann distinguishes between "things" and "devices." A thing in his sense

> is inseparable from its context, namely its world, and from our commerce with the thing and its world, namely, engagement. The experience of a thing is always and also a bodily and social engagement with the thing's world ... Thus a stove used to furnish more than mere warmth. It was a *focus*, a hearth, a place that gathered the work and leisure of a family and gave the house a center. Its coldness marked the morning, and the spreading of its warmth marked the beginning of the day. It assigned to various family members tasks that defined their place in the household ... It provided the entire family a regular and bodily engagement with the rhythm of the seasons that was woven together with the threat of cold and the solace of warmth, the smell of wood smoke, the exertion of sawing and carrying, the teaching of skills, and the fidelity to daily tasks ... Physical engagement is not simply physical contact but the experience of the world through the manifold sensibility of the body. That sensibility is sharpened and strengthened in skill. Skill is intensive and refined world engagement.

Here, in his retrieval of the thing's world and our engagement with the thing, Borgmann has been influenced by Heidegger's fourfold account of things. Obviously, Earth and sky are woven together with mortals. He points out that in

Roman times the hearth was the abode of household gods, though he does not make much of it. Borgmann's account goes beyond Heidegger in emphasizing social and bodily engagement to a degree to which Heidegger seems insensitive. He also steps beyond Heidegger by highlighting the way things focus practices. Practices call for skills and the development of character; the diversity of different characters is joined to each other through participating in a world of practices. In our terms developed earlier, the hearth is the correlational coexistent thing which establishes the world of the household and, correlatively, calls forth its members and calls on their deeper capacities.

Today the hearth, if it exists at all, is no longer the central location in the house although the mantel still remains a place of honor. What has replaced the thing is the "device." The device (the central heating system) provides a commodity, one element of the original thing (warmth alone) and disburdens people of all the elements that compose the world and engaging character of the thing. This world of the thing, its ties to the natural and cultural world and our engagement with that many-dimensional world on bodily, cerebral, and social levels, is taken over by the *machinery* (the central heating plant itself) of the device.

> The machinery makes no demands on our skill, strength, or attention, and it is less demanding the less it makes its presence felt. In the progress of technology, the machinery of the device has therefore a tendency to become concealed or to shrink. Of all the physical properties of a device, those alone are crucial and prominent which constitute the commodity that the device procures.

To make the commodity even more technologically available, the machinery varies radically in the history of technology (wood or coal or oil or electricity or gas). Owing to this radical variability and to this concealment, the machinery becomes necessarily *unfamiliar*. I probably do not know by what means the water is heated in a building. But the device is not just machinery or even most importantly machinery. The device makes available a commodity—warmth. Warmth is what the central heating system is for. Just the opposite of the machinery, the commodity tends to *expand* (become ubiquitous in the house), to remain relatively *fixed* as the means change (from coal to electricity) and to *be familiar*. It follows that—unlike with things—there is a wide division between what a device provides, the commodity, and how it provides this commodity, the machinery. Hence, and this is Borgmann's central insight we saw illustrated earlier with second homes, devices *split* means and ends into mere means and mere ends.

Even though these claims that a thing makes on people are not always experienced as burdensome (as we see from the above account), this very world of the thing and the engagement it calls for can be felt at times as a burden or hassle. The technological device and its refinement *disburdens* people of all these problems by expanding the commodity, so that the world of the thing no longer determines when, in what way, and where it is available. Thus, it disburdens them of the claims that call for engagement. In short, the technological device disburdens people of the thing's world and its claims upon them. The device is considered the more refined the more it lifts these burdens from them. The ideal device is one where, from an experiential standpoint, a commodity can be enjoyed unencumbered by means. A reliable self-regulating central heating system whose maintenance and energy bill are taken care of by a management agency can be taken as a paradigmatic example.

The peculiar way technology dominates things is not limited, of course, to the central heating system. Considering how household technologies have changed, Witold Rybczynkski in *Home: A Short History of An Idea* writes:

> The evolution of domestic technology … demonstrates that the history of physical amenities can be divided into two major phases: all the years leading up to 1890, and the three following decades. If this sounds outlandish, it is worth reminding ourselves that all "modern" devices that contribute to our domestic comfort—central heating, indoor plumbing, running hot and cold water, electric light and power and elevators—were unavailable before 1890, and were well known by 1920. We live, like it or not, on the fair side of a great

> technological divide. As John Lukacs reminds us, although the home of 1930 would be familiar to us, it would have been unrecognizable to the citizen of 1885.

Just as with household technologies, so too with other features of our surroundings and our cultural and natural environment generally. This thing-to-device example is representative of the pattern of the technological transformations of the Earth. Generally then, this transformation is one in which:

> Devices ... dissolve the coherent and engaging character of the pre-technological world of things. In a device, the relatedness of the world is replaced by a machinery, but the machinery is concealed, and the commodities, which are made available by a device, are enjoyed without the encumbrance of or the engagement with a context [that is, the world of the thing].

Borgmann calls this pattern the *device paradigm*. At times I will call it the separation pattern of technology.

We can understand our age as one in which we reduce everything to resources that we want to control. Now we can see that the device pattern is used to get control of these resources. The purpose of the device is to supply people with unencumbered commodities. So now we can develop this picture of our age further. The fuller vision is one in which everything gets reduced to resources, machinery and commodities.

IRONIC CONSEQUENCES

So far we have developed a theory by which we can interpret what has taken and is taking place with regard to the technological transformation in our time. Using this theory we can pass from technological object to technological object, seeing how they more or less fit the pattern. The illustration of the pattern does not commit us as yet to an evaluation of the good or bad of what has taken place. Now we are in a position to begin that task. What are the consequences of this change from things to devices?

Don Ihde finds that technologies transform experience in a "non-neutral" manner. A tool always amplifies in some way certain aspects of normal embodied experience while simultaneously reducing other aspects. A dentist's probe shows the hardness and cavities in a tooth to a degree fingers miss, while the wetness and warmth of the tooth felt by the fingers go undetected by probe. This change Ihde finds is non-neutral because the amplified features are heightened, drawing our attention, while the reduced features tend to go unheeded and are overlooked and often forgotten. Asked what a hearth is for, we find it logical, after having experienced central heating, to answer that it supplies heat, ignoring or not even seeing its other aspects. Extending Ihde's insight makes it more intelligible why we become fascinated with commodities, heedless of what has been reduced. Yet pointing out that this change is non-neutral is not enough. We now want to comprehend what exactly has been hidden from us. We need a language which articulates what is overlooked and forgotten, for then we can see in what ways this change is non-neutral. Our language of things retrieves and focuses this loss. It reveals the general pattern that things are transformed into devices, detaching people from things, their world, and each other.

Ihde further argues that these lost features only tend to recede, thus, implying that they are retrievable. With certain kinds of instruments (not all technological objects are devices, splitting means and ends), this is true. We can easily retrieve the features missed by the dentist's probe. With devices, however, these features do not just tend to withdraw, so that a change of attitude, perception, or act of will could retrieve them. Notice that mere warmth, no matter how expanded the commodity has become, is not a substitute for the thing of the wood-burning stove. Mere warmth could not be the essence of a household; it does not warrant the kind of attention or care, of heeding. Indeed the source itself is concealed and the warmth is suffused throughout the house so that it fails to provide a focus. So warmth is no substitute for the thing because it lacks a world with which to become engaged. More than this, because it is impossible to recover in the mere warmth of the central heating plant,

the full-bodied experience of the hearth, the machinery of devices ineluctably withdraws the world from people. A device is necessarily unfamiliar in the ways that the context of the thing was familiar. Thus the transformation of the thing into a device does not merely tend to obscure possibilities of experience, but its very material structure makes the rich experience of the thing impossible.

Another way of putting this is that devices allow the possibility of only slim points of contact with "narrowly defined aspects of what used to be things of depth." Devices force people to take them as commodity bearers; they leave them no choice. So our way of taking commodities is not a psychological matter, but a real matter. Technology is not only a way of seeing (and for this reason characterizing technology as a vision is perhaps misleading), it is more importantly a way of *shaping*. The very material structure of a device is such that it can be experienced only as holding up a commodity calling for consumption and nothing more.

The implication of this change of shape is alienation. What seemed promising at the outset—relieving people of burdens—leads ironically to disengagement, diversion, distraction, and loneliness. In short, we become not-at-home in the universe. But clearly, simply finding ourselves free from the exclusive use of candles and outhouses does not place us in this alienated position. So how can such positive events as electric lighting and indoor plumbing lead to these ironic results?

To be relieved of famine, cold, darkness, confinement, and other genuine adversities of the human condition was an intelligible and urgent demand for the early phase of modern technology. For the middle-classes of advanced industrialized countries, most, not all, of these kinds of challenges have been met for some time. At the stage of mature technology, the challenges can be quite frivolous. Food processors, electric pencil sharpeners, prepared fishing leaders, automatic cameras, electric knives, and some pain relievers are typical. The basic question here is: Do we need to be relieved of every last and least burden? Aren't some of these burdens actually good in senses that touch our very humanity? When people reflect on these questions they may answer them differently, but when they act, they tend to act in agreement with a vision that seeks to bring everything under control. Ironically, in the wake of such technological success, in the wake of the initial excitement over owning the latest item, the item falls back into the ordinary every day and they become bored. Being bored, they become disengaged and alienated from what may have been a vital practice, such as preparing meals or gardening or photography. Accordingly, they seek diversion. Thus, ironic consequences follow from the disburdenment of every hassle, problem, or felt demand. If we pursue disburdenment in this unchecked and unreflective manner, as people are doing in the stage of mature technology, then these are the results we should expect.

However, it may seem as though we have been just too nostalgic. The disburdenment devices yield "frees us up for other things" as people commonly say. Yet this perspective makes us think that technology is mostly about freedom, as Charles Taylor thinks, when the promise and vision of technology are mostly a promise of happiness. The most unique and devastating critique of technology is not centered on technological freedom, but on the fact that technology *fails most where it succeeds most* at procuring happiness, at procuring the good things of life. As a culture, we think not only that we can use technology to liberate us, but also that we can use it to fill that new possibility space with technologically available goods. In short, what people are freed up for are not other things, but *more commodities*. Then too when people imagine what they are doing as they throw food into the microwave as freeing themselves up for other and more important things, they ignore how pervasive the technological order is. The totality of technological devices is far more consequential than any particular device. The former point can be advanced best by developing the latter first.

Extensively yet unobtrusively this technological way of taking up with the world pervades and informs what people think, say and do. We need an account of technology as *correlational environment*. Organizations, institutions, the ways nature and culture are arranged and accessible all become modeled on the device. As people make more

decisions for consumption against engagement, our average, everyday world is stamped more deeply with the pattern of the device. In other words, devices do not simply liberate people from some things and free them for other and better things. We are surrounded. The things enabling correlational coexistence have nearly disappeared. As the totality of our daily environment changes from an environment of things to an environment of devices, from an environment making demands on people to an environment that is more at their fingertips, this change necessarily entails heedlessness and evokes an attitude of cultural petty homocentrism. So it is important to consider not just the appropriateness of this or that device in a particular context, but to consider what the consequences are of the totality of these devices and people's *typical* use of them. We would expect the consumptive ways of life in such surroundings to be disburdened and disengaged.

So what are people finally freed up for? How do they attempt to use technological means to positively enrich their lives? Typically people use devices to procure entertainment commodities. Hence our culture treats tradition, culture, and nature as resources to be mined. Just as ubiquitously available warmth is not a substitute for the hearth, entertainment commodities are at best insubstantial aspects of the original things. Because they use devices here to procure the delights that matter, final things, these entertainment commodities can be thought of as final commodities. In this respect, to consume a final commodity is no different from consuming an instrumental one that disburdens us of a chore.

THE IRONIC CONSEQUENCES OF FINAL CONSUMPTION

Television is a clear example of a final commodity. Its refinements from the first sets to those of today fit the same pattern as the refinements in central heating and the refinements of devices in accordance with Borgmann's device paradigm in general. So television is an instance of the vision of domination, liberation, and enrichment. It does not make demands on people and is a window of the world, making all the goods of the Earth available, technologically available, to them in their living rooms.

Understandably, television has tremendous appeal to us as a culture. It's where technology comes home to people. The amount of time they spend watching it indicates its power.

> The A.C. Nielsen Company (1989) currently estimates that people in the United States view upwards of 4 hours of television each day. Given the likelihood that such estimates are inflated, let us assume a more conservative estimate of 2 1/2 hours of television viewing per day over the period of a lifetime. Even at this more conservative rate, a typical American would spend more than 7 full years watching television out of the approximately 47 waking years each of us lives by age 70—this assuming an average of 8 hours of sleep per day. Such a figure is even more striking when we consider that Americans have about 5 1/2 hours a day of free time, or approximately 16 years available for leisure of the same 47-year span. From this point of view and based on a conservative estimate, Americans are spending nearly half of their available free time watching television.

Since it is the most popular way people enjoy final consumption, it is worthwhile to examine in detail the *experience* of this form of consumption as we develop the ironic consequences of final consumption. *Television and the Quality of Life* by Robert Kubey and Mihaly Csikszentmihalyi ... do just that, examining systematically the reported experience of television in contrast to the reported experience of other activities people spend time on in their daily life. So we need to look closely at their findings with a view to showing how technology in the form of TV does not fulfill its promise.

Kubey and Csikszentmihalyi find that television is inexpensive and is easily and quickly available for those who have time for it. It helps people to relax and, at times, may help some to retreat before gathering themselves to face a difficulty. People can watch it for prolonged periods without wearing themselves out physically. It tends to bring families together and family members normally do talk with each other while watching. They also feel better than when watching alone.

It is used for news programs, for nature shows, and to present dramas such as *Death of a Salesman*. It could be used to present lectures in chemistry and Plato. It helps connect us with our culture and some of its common stories. On the other hand, "viewing is almost always mildly rewarding in that it provides relaxation, distraction, and escape with minimal effort." It gives people something to do with their time and most report they do want to watch it. With so much of our leisure time taken up with television and with so many benefits, one might mistakenly conclude that people choose to watch because it is better than anything else they could be doing. Yet the actual cumulative benefits they receive from television are rather low and often negative.

Ironically, the reported experience of people viewing television often turns out to be one of disappointment. Not only are chemistry courses not aired because few would watch them, not only do most people gravitate toward watching movies with light and escapist content rather than challenging dramas, but, just as important, half the people who watch television do not use television guides to help them decide what to watch. The stories viewers share, then, are not those shown on public television. The shows tend to support existing beliefs. As Stu Silverman told Kubey, "Television reassures us, it's 'nice,' it doesn't offend or challenge an audience. It is designed to do just the opposite of art, to reassure rather than excite. That often is what people want." Although television does help people to relax, it does not do so any more than other activities such as reading. Moreover, it helps people relax only while watching it and not later as sports and other activities do. Although this study found that television is not a completely passive activity, it is comparatively so. It is not usually challenging, requires little mental alertness, and is reported to exact fewer skills than eating. Only idling was reported to be more passive. Unlike activities that gather and restore a person, a "passive spillover effect" tended to follow watching television, making people feel duller, more passive, and less able to concentrate. Families for which television provides a center also experience this spillover effect carried over into other family activities.

Finally, the positive benefits one receives from television tend to be enjoyed less the more one watches. Heavy viewers are not made happy by watching it; they generally feel worse than light viewers both before and after viewing. Even light viewers do not report themselves to be any happier than average while watching.

Aristotle found amusement, like sleep, to be therapeutic as long as life is oriented around exertion. So, too, Kubey and Csikszentmihalyi find that those who stand to benefit most from television use and need it least. More often television is used to disburden people of problems in ways that do not go to the roots of the problems, are only marginally effective, and, hence, are entirely inappropriate. People disburden themselves of the problem of leisure time that their time-saving and labor-saving devices have created by killing time watching television. They disburden themselves of the problem of loneliness when devices leave them isolated by turning on a device, the television set. Heavy use is higher among singles. Such an answer to loneliness is only a diversion from genuine forms of social engagement. On the other hand, television is often used as a way for family members, usually fathers, to avoid talking with other family members and avoid dealing with family problems. Television resolves the problem of independently ordering one's life, of giving shape to the day. It takes care of boredom. Heavy television viewing is likely driven by a wish to escape, to be disburdened of bad days and bad moods, of personal problems and of alienation from self. Diverting one's attention, it tends to mask the deeper and more real problems a person is having and, hence, leaves these problems unintelligently resolved. Does it meet the task of leading a more rewarding and meaningful life? No. "Happiness is a more complex state than relaxation. It requires a more elusive set of conditions, and is therefore more difficult to obtain." Television seems to "encourage a false sense of well-being in some people," distracting them from and becoming an obstacle to the hard work it takes to realize one's potentials.

More indirectly, we can ask what people are missing when they watch television. When viewers are not pleased with the amount of time they

watch television, the entire reason is not only that it is a low-grade activity, one that many think best fits the phrase "Am I lazy!" Part of the reason, too, has to do with what television is displacing. Many report that they feel as though they should have been doing something else. College educated viewers felt this way more often than other people because "they should have been doing something more productive." Television rearranges life through decreasing the amount of time spent involved with other activities. It at times provides a center for the household, but such a center seems flimsy at best, especially in comparison with other potential centers or centers of the past. In another context, Kubey and Csikszentmihalyi speak of these kinds of centers of life.

> When people are asked what they enjoy most, and enough time is left for a genuine answer to emerge, we often find that the most enjoyable things involve doing something, and usually something rather complex and demanding. Rarely does watching television get mentioned, or any other passive or consummatory activity ... The first reflex for many people is to say that one most enjoys going on vacations, going to movies or restaurants—the typical "leisure" responses in our culture. But as people think more deeply about their real feelings they will mention enjoyable times with their families, and then there is often a point when their faces light up and they say something like: "Actually, the best times in my life have been ..." and start talking with great enthusiasm about designing and sewing quilts, rock climbing, playing music, working on a basement lathe, or about other activities that require concentrated skill, that do not separate the individual from the end result of his or her effort, and that provide the kind of exhilaration and high focused attention of flow. So ... we are still able to keep in sight those *vivid signposts* that show what it is that makes life worth living ... [On] reflecting on such occasions, people often say that not only was the experience enjoyable at the time, but that it helped them grow and become more than they had been. Compared to such optimal experiences, much television watching could be deemed a waste of time ... wasting it amounts to wasting life.

Casting television in terms of the symmetrical

earlier, we can see that the medium is just not enough for humans to make the center of their lives. It does not call forth their humanity in any depth. Hence, Kubey and Csikszentmihalyi worry that by spending so much time viewing television "one may well lose opportunities to grow as a human being."

Kubey and Csikszentmihalyi find that a mistaken cultural assumption underlies much of the appeal of television. For them the mistaken cultural assumption is narrowly one of thinking that physical and mental exertion are bad, and that they are unrelated to human growth and living a worthwhile life. In contrast, for us, the more comprehensive mistaken assumption is that technology generally can fulfill our aspirations for freedom and enrichment. Considered from the standpoint of the vision of technology, television is a paradigmatic example of and not an exception to the unimpeded development of technological culture. As for liberation, it is a commodity which does not make demands—in dress, transportation, or manners, or even having to be at home when a program is aired. Following the device's split between means and ends, people exert themselves in labor and expect to relax completely in leisure. They want amusement, not challenge or disturbance. In terms of prosperity, with video cassette players and hundreds of channels, the most, the most varied, and the latest programs can be watched. Advertisements, too, and the settings of the programs themselves celebrate this prosperity of technology. In short, the incredible attraction of television is that it is the homeplace of the vision people are still spellbound by. It confirms them in that vision and tells them what's what in the universe. Its glamour binds and soothes while simultaneously disappointing them with the flatness and shallowness of its nourishment, its ironic unfulfillment. Television as an exemplar of a final commodity represents the ultimate appeal of the promise of technology. It is, then, the success story of the technology. Television is the vision of technological culture.

So while technology is successful on its own terms when it make goods available, its success is merely a Pyrrhic victory. What makes good things

seduced by a shallow semblance. Thus, technology fails to deliver happiness, not because it fails to make goods available, but because such goods as it does make available turn out to be merely ironic goods. What seemed promising in the appearance is disappointing in the reality. Our aspirations for freedom and happiness go awry when we attempt to procure them with devices.

So does technology deliver the goods? Does technology help people live more rewarding and meaningful lives? As people make the things that count in their lives technologically available, they empty them of depth and they lose them. It is a lesson our culture has yet to learn when we let television turn us into something less than members of the animal kingdom.

1.3.4 ETHICS AND TECHNOLOGY: A PROGRAM FOR FUTURE RESEARCH

DEBORAH G. JOHNSON AND THOMAS M. POWERS

The nineteenth-century essayist Thomas Carlyle declared that "man is a Tool-using animal.... Nowhere do you find him without Tools; without Tools he is nothing, with Tools he is all." We typically think of our tools as inert *things* that are only causally active in the world when put to use by human agents who wield them to accomplish specific purposes. Some authors whose work appears in this book—for example, Richard Sclove and Langdon Winner—have claimed that social and political values are embedded in certain technological artifacts that give them a particular kind of valence when deployed in the ways for which they were designed.

Philosophers of technology Deborah Johnson and Thomas Powers, however, take this line of thinking a step further and argue that "artifacts have intentionality, the property of 'aboutness' or directedness toward the actual world and a future designed world" that is distinct from the intentionality of their designers and users. The intentionality of artifacts is, in their view, built into their functions—that is, the ways in which they typically transform inputs into outputs of various kinds. Although the user's intentions and the context in which the artifact is being used also play a role in determining the outcome of a technologically assisted action, the intentionality embedded in the artifacts themselves is not an eliminatable component of the overall explanation of what takes place. This kind of extended notion of intentionality has been used to describe artificially intelligent computer programs and robots, but these authors think that the notion of surrogate agency, or derived intentionality, can be extended to designed artifacts of all kinds. If they are right, it opens up a rich new philosophical perspective on the designed world that we live in.

Source: From Mitcham (editor). *Encyclopedia of Science Technology and Ethics*, 1E. © 2005 Gale, a part of Cengage Learning, Inc. Reproduced by permission. www.cengage.com/permissions

FOCUS QUESTIONS

1. What do Johnson and Powers mean when they claim that we are living in the anthropocene era?
2. What is a "reason explanation," and why is it necessary for the moral appraisal of action?
3. In what sense can artifacts be said to possess intentionality? How does the intentionality of artifacts differ from that of human agents?
4. What are some examples of how the explanation of a complex action requires reference to the threefold intentionality of designers, artifacts, and users?
5. Does the attribution of intentionality to artifacts entail that designers and users are relieved of moral responsibility? Explain why or why not.

KEYWORDS

anthropogenic, artifactual platform, intentionality, moral agency, surrogate agents, technological function

In this paper we present a program for future study of ethics and technology. Most generally, the analysis involves understanding the role of technology in moral action. On the one hand, technology shapes and is shaped by moral thought and action; on the other, this shaping is rarely acknowledged, let alone understood, by moral philosophers. Thus the program sketched here is aimed at making technology visible as an element of moral philosophy. We lay out a line of reasoning that uncovers the intentionality of the design of technological artifacts, and then we compare human moral action to features of the design and use of technological artifacts. This line of reasoning provides the groundwork for extensive future research. The program description is both a plan of study for our own research as well as a call for other scholars to turn their attention to the issues outlined.

In thinking about the nature of a technology, we argue that traditional philosophical theories of human action and ethics can be usefully extended to technology. Contemporary action theory has suggested a causal model of intentional behavior in humans, and we believe that (with modification) this model is applicable to technology. Indeed, when technology is viewed in relation to a causal model of intentional behavior, the moral nature of technological agency becomes apparent. Similarly, traditional notions from ethics, such as goodness, responsibility, and accountability, can be extended in order to understand technology in a new light.

THE ARTIFACTUAL PLATFORM

The world in which humans act and live is a world filled with human-made objects. In addition to the objects of the natural world, these human-made objects provide an enabling and inhibiting background for human thought and action, and for all of the arrangements of human life. This background influences and informs what we think, how we act, and how we arrange ourselves into units, organizations, and institutions.

By noting the presence of human-made objects, we introduce a distinction between the human-made and the natural world, though we readily admit the two are intertwined. Indeed, they are often so intertwined that it is difficult to separate them. The natural world has been dramatically affected by human activity, and technology is, at least in part, the manipulation of natural potential. Scientific research from the late twentieth and early twenty-first centuries suggests there is very little left of a natural world that is untouched by human agency; the balance, over human history, has clearly shifted toward a relatively larger class of human-made objects. In other words, we are living in the anthropocene, on an increasingly anthropogenic planet (Allenby 2004).

Even though, as a matter of ontology, it will be increasingly difficult to maintain a distinction between the classes of human-made and natural objects, the difference remains significant. The human-made world could be otherwise, and the future human-made world is, to some extent, a matter of human choice and human action. Indeed, work in "normative" design and engineering, seen in the universal design, green engineering, and appropriate technology movements, presupposes that there are morally better (and worse) ways to create the future human-made world. The analysis herein provides these normative enterprises with a philosophical footing.

Moral philosophy has always presumed the natural world as the background for human action and morality, but has failed to recognize the powerful role of the human-made world in moral thought and behavior. Rather than focusing on the background, moral philosophy has concentrated attention on human agency, and the presumption has been that moral action (through human beings) is part of the embodied world. The embodied world has been understood to consist both of natural things and human bodies, though, to be sure, some ethicists have acknowledged that morality might be different if humans had different sorts of bodies or acted in a natural world ordered in a different way. Moral philosophers have considered a typical action to consist of an agent (an embodied being) moving his or her body in some way, even if only in a very small way—a wink, a bit of pressure on a trigger, and so on. If the agent does not move his or her body in some way, then there is no action. Even speech acts require movement of the speech organs, and most philosophers have recognized that humans can commit moral wrongs with mere words.

So our starting place is the idea that human agency operates in an embodied world, noting how the embodied world includes both human-made and natural objects. But we want to call attention to the *normative* features of the human-made part and come to grips with the moral importance of technology in constituting the background for human action. We will call the human-made part of the embodied world, as far as it concerns human action, the *artifactual platform*. This platform is the class of constructed objects and systems of objects that are created by and come to influence human action.

Often, descriptions of action incorporate human-made objects into the action. For instance, when we say "John shot Bill," use of a gun is implicit; when we say "Mary flew to London," use of an airplane is presumed; and so on. This feature of descriptive language is what Joel Feinberg (1970) has called the "accordion effect." We can choose an expanded description that includes the artifact, or a collapsed version that conceals it.

When those who study action from the normative point of view use narrow or collapsed descriptions, the technological component is glossed over. What is missed is that particular movements of an agent's body could not have had their associated effects were it not for an artifact. Noting the artifacts involved in moral behavior is the first step in gaining a better understanding of the role of the artifactual platform in morality. Becoming aware of this platform allows us to see that a good deal of moral behavior is action *with* technology. In this respect, moral actions, agents, and patients are not sufficient for an ontology of morality; artifacts are also part of the moral world. The task of understanding the role of artifacts in morality is, then, a matter of recognizing the difference it makes for humans to live in a world with the particular artifacts that currently exist or might exist in the future.

Nevertheless, realizing that moral action takes place *with technology*, and on or from an artifactual platform, does not go far enough. As indicated, technological artifacts with their particular features are matters of human choice. Just as humans deliberate about and choose their actions, some humans (artisans and engineers) deliberate about and create artifacts; other humans (consumers and users) choose and employ artifacts that enable and constrain moral action. Human agency is significantly affected by technological artifacts. It may be augmented, constrained, or merely altered. The design, availability, and employment of technology shapes what humans can do, and what they end up doing.

What, then, is the significance of technology? Technology expands and constrains the range of

human moral behavior, and changes the character of that behavior. Technology is far from neutral in its combination with human behavior. Can one say that it has moral agency? This question can be pursued by considering relations between human moral agency and technology.

THE MORAL AGENCY OF TECHNOLOGY

The question of the moral agency of technology can be used as an entry point for exploring the role of technology in morality. Grounding it in philosophical concepts, the analysis starts with the traditions of ethical theory and action theory and the accounts of human moral agency they provide. In ethical theory, the standard account of the responsibility of moral persons (acting without technology) says that individuals are primarily responsible for their voluntary, intended behaviors. In action theory, there is a broader account of intentionality, in which intentional states ("intendings" as well as desires, beliefs, plans, etc.) are the causes of action. The intentionality of these states is a property that relates them to states of affairs and objects in the actual world and in possible worlds. Intentionality, then, is "aboutness" or directedness. On this view, voluntary action or intended behavior is understood to be outward behavior caused by a complex of internal mental states. By stipulating the specific kind of intending, desiring, and believing that causes a particular action, philosophers have distinguished moral action from nonmoral behavior. Because the outward behavior in moral action is the result of these internal mental states, it is amenable both to a causal explanation and to a "reason explanation" (see Davidson 2001). That is, when we ask why someone acted in a particular way, he or she can offer antecedent intendings, beliefs, desires, and other intentional states as reasons for the action.

The standard philosophical account is spelled out in contemporary work in ethical theory and action theory, but the roots of the account are much older. The subject matter of moral appraisal even as far back as Aristotle (384–322 BCE) has been understood to be intended, voluntary behavior. This is action, conduct, or the commission of a deed, as opposed to "mere" reaction or nonvoluntary behavior. In contemporary action theory, Aristotle's basic view is elaborated upon, and this produces the following conditions for moral action. First, there is a potential agent with an internal state. The internal state consists of intentional mental states, one of which is, necessarily, an intending to act. Together, the intentional states (e.g., belief that X is possible, desire to X, plus an intending to X) constitute a reason for X-ing. Second, there is an outward, embodied event—the agent does something, moves his or her body in some way. Third, the internal state is the cause of the outward event; that is, the movement of the body is rationally directed and is an action insofar as it is caused by an internal state. Fourth, the outward action has an outward effect. Finally, the effect has to be on a patient—the recipient of an action that can be harmed or helped. Moral patients are typically human beings, but the class may include other beings or things as well. Some ethicists now include higher functioning animals, entire species, and even ecosystems in the class of moral patients, and clearly technology does seriously affect ecosystems and nonhuman animals.

The convergence of these parts of ethical theory and action theory has produced a plausible account of the connection between thought and action and has helped locate the focal point of moral agency. We adopt this account as the framework in which to consider the moral agency of technology. In other words, whether or not or in what ways technology has moral agency can best be revealed by comparing features of technology with the standard account of moral action as derived from ethical theory and action theory.

Interesting work has been done in the late twentieth and early twenty-first centuries along these lines, as philosophers have turned to consider the possibility of nonhuman moral agents (Allen, Varner, and Zinser 2000, Floridi and Sanders 2001, Brooks 2002, Kurzweil 1999, Danielson 1992). Most attention has been given to artificially intelligent computers as the best candidates for agency. Computers have drawn attention

of intelligence. Some philosophers of artificial intelligence (AI) seem to think that intelligence can emerge out of the complex states of computers. This view implies that the ability of a computer to generate intentional states on its own would go a long way toward making it like a human moral agent. (Researchers in AI are primarily interested in engineering robotic computers to do things such as sense, recognize, navigate, and modify, and not, in the main, concerned with the deeper implications of AI for a philosophical account of intelligence.) A thrust of the account here is to draw attention away from the project of considering intelligence and computers, and instead to explore technological artifacts more broadly, as entities that have intentional states that are not mental states.

At the heart of our argument for the moral significance of technology is the claim that artifacts have intentionality, the property of "aboutness" or directedness toward the actual world and a future designed world. One of the reasons so little attention has been given to ethics and technology seems to be a failure to recognize the intentionality designed into technological artifacts. On the one hand, the only type of intentionality of interest to ethicists has been the type found in the mental states of *human* agents. With its focus on human agents, ethical theory has not recognized the importance and relevance of the design and use of technological artifacts by human agents. On the other hand, scholars in science and technology studies have introduced the idea of technology having a kind of agency (Law 1987, Callon 1986). However, they have not recognized the ethical implications of this move. Nor have they related technological agency to the broader philosophical literature on action. The argument in this essay brings ethical theory and action theory to bear on the moral agency of technology.

Because the program outlined here builds on our claim that artifacts have intentionality, it will be helpful to discuss the theoretical apparatus traditionally used to describe intentionality in moral action. In order for a human action to be both open to "reason explanation" and subject to

collection of intentional mental states connected to the action in some fairly specific ways. Agents are subject to moral appraisal in virtue of those intentional acts that have morally relevant effects on moral patients. Intentional acts are caused by a variety of intentional states and/or entities: beliefs, desires, intendings, maxims, plans, and the like. An agent is a being who acts, with the cause of the action originating in the agent's mind as the complex of intentional states. The cause of the action is the primary reason for the action, and the cause as a whole can be seen as a collection of intentional states that serve as a "reason explanation" of the action. Intentional entities are entities that are capable of having intentional states; intentional actions are those actions that are caused by intentional states.

Our extension of this view of agency does not entail that artifacts have mental states or the ability to intend. We claim only that artifacts have intentionality or directedness at users and environments, and that this intentionality is causally efficacious. In proposing that intentionality is designed into technological artifacts, we avail ourselves of a quite general definition of intentionality, according to which it is the property of something, such that it is directed at or represents an object or state of affairs. The term *intentionality* is broadly construed so that intentional entities can be states of mind, sentences, speech acts, maps, or any designed object. Though this view of intentionality is quite broad, we nonetheless agree with the traditional view that humans are intentionality-generating beings. Their states of mind are directed at or about objects and states of affairs, and it is this original power of mind as intentionality generating that accounts for the intentionality in nonmental entities.

Humans have the ability to externalize their intentional states in speaking and writing. Spoken and written declarative sentences are intentional, just as are the beliefs that they express. While sentences and signs originate in the processes of the mental realm, these entities come into being only when they are expressed outwardly. Clearly, some intentional entities remain *internal* to humans, such as mental states of belief, desire, and visual

actions of human moral agents in that the intentional entities cause the actions and count as reasons why the agent committed the act. As for the external intentional entities, once they come into being and are (by definition) physically separated from the human who generated them, they still rely on a community of intentionality-generating beings (interpreters) in order to be intentional—in order for their intentionality to be grasped. Examples are maps, chairs, sentences in a natural language, and works of art. External intentional entities, like their internal counterparts, can cause and explain action. For example, the stop sign causes drivers to step on the brakes and bring their vehicles to a stop; the speech act of commanding individuals to behave in a certain way may cause individuals to do what is commanded; and so on.

The internal/external distinction in intentional entities takes into consideration the kinds of intentionality in human minds, in tangible expressions such as sentences and speech acts, and in representational states that are found in designed artifacts. Internal intentional states are those that necessarily remain mental; external intentional states, by contrast, are expressed in the form of entities that exist outside of the mind. An internal intentional state such as a belief often leads to an external intentional entity by means of a process not yet fully understood, but still assumed to be causal in nature. We argue that designed artifacts such as maps, computer programs, cars, and the like are externalized expressions of internal intentional states. They are intentional entities that cause action with morally relevant effects.

The most difficult part of the account here is the claim that things other than mental states can be about, be directed at, or represent objects and states of affairs. This claim seems noncontroversial when applied to sentences, speech acts, and maps. For instance, John R. Searle (2001) describes maps and house blueprints as intentional entities. Thus it should not be controversial when it comes to technological artifacts. While we claim that technological artifacts are intentional entities, we acknowledge that in the standard account of agency and action, agents have a specific intentional state of intending to perform a particular action, plus some more basic intentional states such as beliefs and desires. Because we claim that artifacts are intentional entities, the obvious question is what kind of intentionality do they have? That is, do they have something akin to the basic intentional states of humans, such as beliefs and desires, or something like the specific states of intending?

THE FUNCTIONALITY AND INTENTIONALITY OF ARTIFACTS

Our argument for the intentionality of technological artifacts is based on a particular understanding of the intentional states that artifacts can have. These intentional states cannot be fully understood without reference to the functions of the artifact. Accordingly, our account of the functionality of artifacts will be developed by answering three questions. What are functions in an artifact? How do they get into the artifact? What do users do with functions?

What Are the Functions in an Artifact?

Typically artifacts are thought to have functions, and their functionality is framed in terms of purposive or teleological explanation. While we do not reject this approach, we want to suggest a different view—one that allows for the flexibility we find in the design and use of artifacts. We base our understanding of the functionality of artifacts on the model of mathematical functions. An artifact has a function when it takes some input from a domain of human behaviors and produces a result within a range—what we generically call the output. The behavior of the user with the artifact fits the mathematical model of functions in that it consists of a relational triple: input, rule of transformation, and output. In the case of both mathematical functions and artifacts, one of two things can happen in the functional transformation. Either an input maps onto exactly one output (in which case the relation is one-to-one) or many different inputs map onto one output (a many-to-one relation). The definition of

a function precludes the possibility that a particular input will deliver varying outputs (except in the case of artifacts such as slot machines whose one output is to produce varying outputs). This is an important condition for mathematical functions as well as artifactual ones. An artifact ceases to be useful (or even sometimes safe) when its output is unpredictable (except, again, when unpredictability is the designed output), and this is exactly what happens when a user gets different outputs for the exact same input on different occasions.

Here is an example of a technological function. A designer of a braking system for cars would model input by considering reaction times, leg position, pedal pressure, and stopping force for drivers who wish to control a typical car by pressing on the brake pedal. This process of design begins to reveal how the artifact becomes intentional; the input model is "about" driver capabilities and driving conditions—what we can gloss as "input" and "environment" aspects of the model. The transformation rule for the function, which is embodied in the mechanical parts of the braking system, turns those anticipated inputs into a result: The car slows at an appropriate speed. This is how the intentional states are actually manifested in the artifact; they are "materialized" in the way the artifact transforms the input. A successful braking system will incorporate realistic reaction times and pressures for the vast majority of drivers, and will reliably transform those inputs into the safe braking of a car under most conditions. A proper braking system will not map the different outcomes "stop the car" and "accelerate the car" to the exact same driver behavior. Design functions, like mathematical functions, are not one-to-many relations.

When an artifact *appears* to function differently with the same inputs, either the artifact is broken or there is a mistake about the sameness of inputs. The input mode for many complex artifacts such as computers is context dependent. For example, when the input of "striking the return key" on the keyboard yields different results at different times, this is because the computer is in different states during the respective inputs. In some programs, a query can be answered affirmatively by striking the return key. In others—word processors, for example—striking the return key places a hard return in a document. The lesson is that inputs are always tied to context. The condition that the artifactual functions borrow from mathematical ones reveals that there will never be more than one output for an *input in a context*. We may get spaces in some word-processing documents when we push the return key, and affirmations to queries when running other programs, but we will never get spaces sometimes and affirmations other times, in the exact same input context.

How Do Functions Get into an Artifact?

Crucial to this account is the fact that transformation rules of functions cannot be built into artifacts without applying intentional models of users and the world in which they operate.

There are two immediate senses in which the intentionality that begins with design is connected to technological artifacts in use. The act of design always requires intentionality—the ability of a designer to represent, model, perceive, and the like. Similarly, the use of an artifact—grasping a tool, following the user's guide—requires typical forms of cognition that feature intentionality. But there are deeper ways intentionality connects to designed functions and uses, ways that go beyond the intentionality of designers and users. When designers design artifacts, they poise them to behave in certain ways. Those artifacts *remain* poised to behave in those ways. They are designed to produce unique outputs when they receive inputs. They are directed at states of affairs in the world and will produce other states of affairs in the world when used. The telephone is "about" typical human fingers and ears, auditory capacities, and the physics of sound—it is intentional with respect to certain organisms and their environments. In a complicated way, the intentionality of the telephone is required to make it work as a communication device. But the telephone is also directed at certain social facts; it is about a world in which individuals want to talk with others who are beyond the reach of (unassisted) human

voices. The telephone also requires that users memorize or keep a record of numbers attached to persons. Otherwise, a potential caller will not be able to use the telephone. Long after the designer has poised the artifact, the functions still reside in it and make complex actions possible. The argument here receives support from an analysis by Fred Dretske (1989) of what he terms the "design problem," as exemplified by how to get a mechanical system to do something that its designers find important, such as how to get a temperature indicator to be a switch for turning on a furnace.

What Do Users Do with Functions?

Users do not merely comply with the behavioral requirements designed into artifacts; they do not merely "satisfy" the model of use. They can add to the functions of an artifact by envisioning an unanticipated input that yields a novel output. This envisioning itself begins as an intentional state in the user, but it is then manifest in outward ways. An example of this is when someone picks up a television and throws it at an attacker to stop the attack. Here the user sees that by providing a particular kind of input (lifting and throwing), the television can be used to produce an output that it was not originally designed to produce.

The intentional states of artifacts are the result of the work of the artifact designer; designers *mold* intentionality into artifacts by concretizing the intentional models so that they enable the transformations promised by the functions. Users then deploy these functions by supplying inputs to the artifacts, under the prescribed conditions. Our argument is thus more than that the intentionality of designers and users becomes operative when artifacts get put to use. Our claim is that artifacts are in some sense chunks of intentionality, externalized by artifact designers and deployed by users in particular contexts.

When the intentionality and functionality of artifacts are seen in this light, it becomes difficult to locate precisely the agency in human actions with technological artifacts. There is intentionality in the mind of the artifact user, in the intentional states and functions of the artifact, and in the designer who created the intentionality and functionality embodied in the artifact. What may begin as the intentional model of a designer gets molded into an artifact and then deployed by the user. Hence, there is a complex of agency with human and nonhuman components.

We thus acquire a picture of moral action with technology as a complex combination of the intentionality of artifact designer, the intentionality of the artifact, and the intentionality of the user. Does this mean that artifacts are moral agents? If we return to the standard account of moral agency, it is now clear that artifacts meet most but not all of the conditions. Remember that on the standard account, human moral agency includes the stipulation of a potential agent with internal mental states, and one of these states is an intending to act. The agent does something, moves his or her body in some way, such that the internal states are the cause of the movement. The internal, mental states are thus also the reason for the action. The movement or behavior has an effect on a moral patient, someone or something that can be harmed (or helped).

Our analysis of human-action-with-artifact overlaps significantly with standard (nontechnological) human action, even though it locates agency in the triad of designer, artifact, and user. We have found that intentional states are spread out over designers, artifacts, and users, so that the action of the human-agent-with-artifact is caused by intentional states in each member of the triad. A complete reason explanation must include an account of the intentional states and functions of the artifact, because these states and functions play a causal role in the eventual action. The causal role of the artifact is necessary, but not sufficient, for the effect on the moral patient. True, artifacts alone are not agents, nor are their intentional states in any way internal mental states. Likewise, artifacts alone do not intend. But the intentional states of artifacts shape and cause external or embodied movement, both in terms of functional inputs of users and in terms of artifactual output. And intentional, caused, embodied movement can have morally relevant effects on patients. Thus, the intentionality and functionality of artifacts

are important components of a full picture of moral action.

This account has implications for the notion of moral responsibility. Because philosophers and others may resist the idea of any kind of agency or even intentionality being attributed to technology because it may appear to deflect responsibility from human actors, it is appropriate to consider the issue of responsibility in a case study. Can technological artifacts be said to bear moral responsibility, or even to be morally good or bad entities?

AN ILLUSTRATION: THE MORAL EVALUATION OF COMPUTERS

At first glance, the idea of artifacts bearing moral responsibility appears implausible. There is, however, a form of human moral responsibility that is applicable to certain kinds of computer systems that may have broader application to other technologies. We refer here to the responsibility of human surrogate agents to their clients. Human surrogate agents are those who act on behalf of others. For example, lawyers, tax accountants, estate executors, and managers of performers and entertainers pursue the interests of their clients. The behavior of these agents is evaluated in terms of how well they pursue their client's interest while staying within the constraints and expectations associated with their roles. Like surrogate agents, computer systems pursue interests of their users; hence, their behavior can be evaluated in terms of how well they pursue the interests of their users.

If computer systems can be understood as surrogate agents for their human users, it would seem that role morality can be extended to computer systems, and this is a reason for attributing moral responsibility to computer systems and for morally evaluating such systems. In essence, the suggestion here is that the concept of role morality can be understood as a set of constraints on behavior, based on the interests of others, and can be applied to the functionality of particular computer systems. Just as human surrogate agents are evaluated in terms of whether they adequately their clients, one can evaluate computer systems in terms of how they represent and pursue the user's interests. Such an evaluation would involve many aspects of the system, including what it allows as user input and how it goes about implementing the interests of the user.

Consider the search engine surrogate that pursues a user's interest in finding web sites on a particular topic. Whether the search engine lists web sites in an order that reflects highest use, or fails to list some sites, or gives priority to sites for which the owner has paid to be listed—all of this can have moral implications (Introna and Nissenbaum 2000). We might say, then, that the computer system takes on a third-person, interested perspective, either of the user or of someone else. Several important questions arise. Does the system act on the actual user's interests, or on a restricted conception of the user's interests? Does the system competently pursue the user's interests, without pursuing other, possibly illegitimate interests such as those of advertisers, computer hardware or software manufacturers, government spying agencies, and the like? Are faulty or buggy computer systems analogous to misbehaving human surrogate agents? Do they fail to do the tasks (or to adequately do the tasks) that users employ them to do?

The foregoing suggests the kind of moral evaluation that can be made when computer systems are seen as surrogate agents. Tax preparation programs perform like tax advisers; contract-writing programs perform some of the tasks of attorneys; Internet search engines seek and deliver information like information researchers or librarians. Other types of programs and computer systems serve the interests of users, but there are no corresponding human surrogate agents with whom to compare them. Spyware programs uncover breeches in computer security, but when they do so for the user, they do not replace the tasks of a private detective or security analyst. Increasingly, computer systems do more for us than human surrogates could do. This is why it is all the more important to have a framework for morally evaluating computer systems, especially a framework that acknowledges that com-

pursuing the interests of their users and can misbehave in their work on behalf of users.

To claim that computer systems (and possibly other technologies) have moral responsibility and can be morally evaluated is *not* to claim that the responsibility or blameworthiness of users or system designers is thereby diminished. We anticipate that the standard response to our argument will be that the attribution of responsibility to various agents is a zero-sum situation—that designers are "let off the hook" when we turn to the moral evaluation of computer systems. In response, we deny that moral evaluation is zero sum. Computer systems behave. Their behavior is intentional, and it can have effects on humans and can be morally appraised independently of an appraisal of their designers' behavior. What the designer does and what the computer does (in a particular context) are different, albeit closely related. To think that only human designers are subject to morality is to fail to recognize that technology has intentionality, and its intentionality plays a causal role in the effects that computer systems can have on moral patients.

So the point of emphasizing the moral responsibility and moral evaluation of computer systems is not to deflect responsibility away from system designers or users. Because a computer system is conceptually distinct from the computer system designer and user, all three should come in for moral scrutiny. Computer systems are an interesting case here because they are becoming increasingly sophisticated, in both technical and social dimensions. Though the first computer systems may have been simple utilities or "dumb" technologies designed to help humans connect phone calls, calculate bomb trajectories, and do arithmetic, computer systems are increasingly taking over roles once occupied by human surrogate agents. This continuous change would suggest that, somewhere along the way, computer systems changed from mere tool to component of a complex agent. Now, it can no longer be denied that computer systems have displaced humans—both in the manufacturing workforce, as has long been acknowledged, and more recently in the service industry. It would be peculiar, then, for users to recognize that computers have replaced human service workers who have always been supposed to have moral constraints on their behavior, but to avoid the ascription of similar moral constraints to computer systems.

We introduced this discussion of computer systems as a way of opening up the possibility of technology bearing moral responsibility and being subject to moral evaluation. The challenge of the program we propose is to explore this territory in relation to both smart as well as more mundane (less complicated) technologies. The larger program will have to come to grips with the triad involved in moral action and agency: designers, artifacts, and users.

CONCLUSION

The line of reasoning developed here sketches an account of the role of technology in moral action. We began with the distinction between natural and human-made objects and noted that moral philosophy has neglected the importance of the artifactual platform in which human action occurs. We argued that artifacts have intentionality and gave an account of this intentionality using the functionality of artifacts and their directedness at states of affairs in the world; in this way, artifacts are comparable to speech acts. Building on our account of the intentionality of artifacts, we considered whether artifacts have moral agency. Here we argued that there are three forms of intentionality at work in moral action with technology: the intentionality of the artifact designer, the intentionality of the artifact, and the intentionality of the artifact user. Allowing for the agency of artifacts does not diminish the responsibility of human actors. To address the issue of the responsibility and moral evaluation of artifacts, we examined computer systems as surrogate agents. We argued that the responsibility of human surrogate agents provides a good model for making sense of the responsibility of computer systems. Computer systems can be morally evaluated in terms of their roles in relation to users. We have long known that computer systems can err; our account suggests that they can also misbehave.

The set of issues discussed here constitute a program for future research. Technology has not

been a significant focus in moral philosophy, and yet it shapes the human moral universe in significant ways. Attention to technology promises to open up a range of interesting, complex, and important philosophical issues.

BIBLIOGRAPHY

Allen, Colin; Gary Varner; and Jason Zinser. (2000). "Prolegomena to Any Future Artificial Moral Agent." *Journal of Experimental and Theoretical Artificial Intelligence* 12(3): 251–261.

Allenby, Braden R. (2004). "Engineering and Ethics for an Anthropogenic Planet." In *Emerging Technologies and Ethical Issues in Engineering*. Washington DC: National Academies Press.

Brooks, Rodney A. (2002). *Flesh and Machines: How Robots Will Change Us*. New York: Pantheon.

Callon, Michel. (1986). "Some Elements of a Sociology of Translation: Domestication of the Scallops and the Fishermen of St. Brieuc Bay." In *Power, Action, and Belief: A New Sociology of Knowledge?* ed. John Law. London: Routledge and Kegan Paul.

Danielson, Peter. (1992). *Artificial Morality: Virtuous Robots for Virtual Games*. London Routledge.

Davidson, Donald. (2001). *Essays on Actions and Events*, 2nd edition. Oxford: Clarendon Press.

Dretske, Fred. (1989). "Reasons and Causes." *Philosophical Perspectives* 3: 1–15.

Feinberg, Joel. (1970). "Action and Responsibility" Chap 6 in *Doing and Deserving Essays in the Theory of Responsibility*. Princeton, NJ: Princeton University Press.

Floridi, Luciano, and J. W. Sanders. (2001). "Artificial Evil and the Foundation of Computer Ethics." *Ethics and Information Technology* 3(1) 55–66.

Introna, Lucas D., and Helen Nissenbaum. (2000). "Shaping the Web: Why the Politics of Search Engines Matters." *Information Society* 16(3): 169–185.

Johnson, Deborah G., and Thomas M. Powers. (2005). "Computers as Surrogate Agents." In *Information Technology and Moral Philosophy*, eds. Jeroen Van Den Hoven and John Weckert. Cambridge: Cambridge University Press.

Kurzweil, Ray. (1999). *The Age of Spiritual Machines When Computers Exceed Human Intelligence*. New York: Viking.

Law, John. (1987). "Technology and Heterogeneous Engineering: The Case of Portuguese Expansion." In *The Social Construction of Technological Systems*, eds. Wiebe E. Bijker, Thomas P. Hughes, and Trevor J. Pinch. Cambridge, MA: MIT Press.

Searle, John R. (2001). *Rationality in Action*. Cambridge, MA: MIT Press.

PART TWO

Contemporary Technology and the Future

2.1 Security and Surveillance

2.2 Artificial Intelligence and Robotics

2.3 Nanotechnology

2.4 Internet and Social Media

2.5 Biotechnology

2.6 Energy and the Environment

2.1 SECURITY AND SURVEILLANCE

2.1.1 THE CONSEQUENCES OF THE INFORMATION REVOLUTION

MAX BOOT

Most of us are familiar with the consequences of the information revolution on civilian life—personal computers, the Internet, blogs, online shopping, and so forth—but unless one is a member of the military, one probably has only a hazy idea of the effect of information and communications technologies (ICT) on the battlefield. Many authors have talked about a "revolution in military affairs" brought about by the integration of advanced ICT capabilities into the arsenals of modern armies, but few have studied its effects and understood its consequences and limitations as thoroughly as Max Boot, a defense analyst and military historian. In this selection taken from his book *War Made New: Technology, Warfare, and the Course of History—1500 to Today*, Boot provides a glimpse into the awesome powers of the U.S. military machine and, in particular, the additional capabilities for land, sea, and air warfare made possible by the use of advanced ICT. But the same information technologies that the U.S. employs can also be used by terrorists to wreak havoc on the twenty-first-century global information economy. Boot suggests that to counter these new threats security analysts need to think beyond "technological fixes" toward a paradigm shift in the way in which they prepare for and wage war.

FOCUS QUESTIONS

1. How have the advanced communication technologies employed by U.S. land forces allowed them to overpower more numerous adversaries? What is the "Achilles heel" of these technologies in counterinsurgency operations?
2. Why is the advantage that the U.S. military has in space slowly eroding? What can potential adversaries do to shield themselves from satellite surveillance?
3. In what ways does U.S. reliance on high-tech systems create vulnerabilities of its own? What kind of countermeasures does Boot think are needed to meet these new kinds of threats?
4. What is meant by the "tooth to tail" ratio, and what needs to be changed in the way in which military bureaucracies operate in order to reduce these ratios?

KEYWORDS

FORCEnet, Future Combat System, legacy systems, smart bombs, space warfare, the fog of war, visual stealth technology

Source: "The Consequences of the Information Revolution" from *War Made New* by Max Boot. Copyright

It is tempting to argue, based on the difficulties that U.S. forces have encountered in Iraq (and the similar problems of the Israeli Defense Forces in Lebanon, the West Bank, and Gaza Strip), that the effect of the Information Revolution on warfare has been overstated by some enthusiasts. There is some truth to this—there *has* been an awful lot of hype—but it would be a mistake to go too far in dismissing the results of recent advances. That would be akin to denying the effect of the Industrial Revolution in the early 1900s simply because the British had a lot of trouble subduing the Boers. Improvements in technology have always had a more immediate impact on conventional than on irregular warfare. But that doesn't mean that the art of war has remained static. Indeed the tactics of Islamist guerrillas, who rely so heavily on the Internet, cell phones, and satellite television—all of which barely existed in 1980—show just how much things have changed.

Not all of the changes wrought by the Information Age are obvious at first glance, because the basic military systems of the early twenty-first century look roughly similar to their predecessors of the Second Industrial Age. Military analyst Michael O'Hanlon notes that "basic propulsion systems and designs for aircraft, ships, and internal-combustion vehicles are changing much more gradually than in the early twentieth-century, when two of those three technologies had only recently been invented." The average speed of a U.S. Navy destroyer, for instance, has not increased in the past one hundred years. The U.S. Air Force continues to rely on B-52H bombers last built in 1962. And the Marine Corps still uses helicopters that flew in the Vietnam War. What has been changing with great rapidity since the mid-1970s is the communications, targeting, surveillance, and ordnance technology that can make such "legacy" systems considerably more potent.

Land Warfare

Advanced armies are still structured, as they have been since the 1940s, around armored forces complemented by infantry troops who move by armored vehicle, truck, and aircraft. The best tank in the world is probably the American Abrams (of which the U.S. has nine thousand), but the British Challenger II, the German Leopard II, the Israeli Merkava Mk. 4, and the Russian T-80 and T-90 come within striking distance. All modern tanks have stabilized turrets, night-vision capabilities, laser range-finders, and targeting computers that allow them to fight in conditions—on the move or in the dark—that would have stymied earlier models. In addition, composite or reactive armor offers far more protection than in years past, and main guns firing depleted uranium rounds have far more penetrating power. Armored personnel carriers and infantry fighting vehicles, such as the American Bradley Fighting Vehicle and Stryker and the Russian BMP and BTR, are essentially light tanks, some running on wheels, others on tracks, that sacrifice armor and armaments for extra room to carry infantry, command-and-control suites, or other cargo. Self-propelled artillery and rocket systems are also mounted on armored chassis.

Armored vehicles have improved over the years. But so have antiarmor weapons. These range from heavy missiles such as the U.S. Hellfire and Russian Ataka-V fired from vehicles or aircraft to handheld versions such as the U.S. Javelin, the Franco-German Milan, and the Russian Kornet. In addition, even the most advanced tanks can be disabled by other tanks, massive mines, aerial bombs, or artillery shells. The full impact of advances in antiarmor technology has not yet become apparent, because the forces that have fought modern tanks in recent years—Iraqis, Palestinians, Chechens—have not possessed the latest defensive weapons. But the U.S. success in wiping out Iraqi tanks from standoff ranges suggests that, in the constant struggle between offense and defense, the advantage may have shifted against heavy armor.

The U.S. Army is responding to these changes by budgeting at least $124 billion to develop a Future Combat System that is supposed to replace much of its current armored force with a family of lighter vehicles, manned and unmanned, that will have stealth designs that will make them harder to detect and hybrid-electric engines that will lessen their fuel requirements—

Abrams, because it increases demands on vulnerable supply lines. Future vehicles will feature advanced composite armor designed to deliver more protection than current models for the same amount of weight, but they will rely for protection less on armor and more on locating and destroying the enemy before they are attacked. Critics believe this is placing too much faith in "perfect situational awareness," and that these vehicles will not be of much use against guerrillas who can strike with no warning.

As usual, the infantryman's tools have changed least of all. A modern soldier has better protection than his forefathers if he wears Kevlar body armor, but his firepower—which comes primarily from a handheld assault rifle like the M-16 or AK-47 and from a variety of crew-served mortars and machine guns—does not vary significantly from that of a G.I. in World War II. A replacement for the M-16 known as the XM29 is under development but it is hardly revolutionary. In addition to shooting the same 5.56 mm rounds as the M-16 out of one barrel, it will have another barrel that can fire 20 mm high-explosive airburst projectiles to a range of half a mile. These mini grenades will come with embedded microchips that will control when they explode, allowing them to kill enemy fighters who might be lying flat on the ground or hiding behind a berm. Alternatively, nonlethal projectiles like rubber balls could be substituted for crowd-control situations. This is not terribly different from the capability afforded by grenade launchers attached to today's M-16s. Electronic guns that are capable of spitting out a million rounds a minute have also been developed. They might permit a soldier to stop an incoming rocket-propelled grenade with a solid wall of lead. But such weapons are years away from being fielded.

Unfortunately for Western infantrymen, the proliferation of small arms can put low-tech foes on an almost equal footing with the representatives of the most advanced militaries. There are 250 million military and police small arms knocking around the world, and more are being manufactured all the time by at least 1,249 suppliers in ninety countries.

The salvation of Information Age infantry, at least when they are conducting conventional operations, is their ability to use a wireless communications device to call in supporting fire on exact coordinates. It is doubtful that any military force will again enjoy the preponderance of power of a Kitchener at Khartoum, but Americans dropping JDAMs on Afghan tribesmen armed with Kalashnikovs—or even on Iraqi soldiers with outdated T-72 tanks—came close. The American edge decreases considerably, however, when its troops have to deploy for peacekeeping or counterinsurgency operations that leave them exposed to low-tech ambushes.

Naval Warfare

Navies remain divided, as they have been since the dawn of the Second Industrial Age, into aircraft carriers, submarines, and surface ships. The major difference is that blue-water naval competition has disappeared after more than five hundred years. No one even tries to challenge the U.S. Navy anymore on the high seas. Virtually every other navy in the world is little more than a coastal patrol force.

The U.S. has twelve aircraft carriers, nine of them *Nimitz-class*, nuclear-powered supercarriers that can carry more than seventy high-performance aircraft such as the F/A-18E/F Super Hornet. A tenth supercarrier is in the works. No one else has a single one. France has the world's only other nuclear-powered aircraft carrier, the *Charles de Gaulle*, but it's half the size of the *Nimitz*. Russia has one aircraft carrier, the *Admiral Kuznetsov*, that rarely leaves port, and it has sold another one, the *Admiral Gorshkov*, to India. Britain has three small *Invincible-class* aircraft carriers that are used only for helicopters and vertical-takeoff Harrier jets. France, Italy, Spain, Japan, and South Korea have similar helicopter carriers in the works. These ships are comparable to the U.S. Navy's twelve amphibious assault ships, which transport helicopters, jump jets, and marines.

Whenever they leave port, U.S. capital ships are surrounded by surface and submarine escorts. Twenty-four *Ticonderoga-class* cruisers and forty-five (and counting) *Arleigh Burke-class* destroyers

come equipped with Aegis phased-array radar that can track up to nine hundred targets in a three-hundred-mile radius. These surface combatants can also operate on their own or in conjunction with smaller vessels such as frigates and minesweepers.

In World War II, ships that didn't carry aircraft were limited to firing torpedoes or heavy guns with a range of less than thirty miles. Starting in the 1960s some submarines were equipped with intercontinental range ballistic missiles, but their targeting was so imprecise that it made no sense to equip them with conventional warheads. Ballistic-missile subs became a mainstay of nuclear deterrence. The development of accurate cruise missiles starting in the 1970s allowed submarines and surface combatants to hit land targets hundreds of miles away with conventional ordnance. Improvements in torpedo design, including the development of rocket-propelled supercavitating torpedoes, also allow submarines to do more damage in their traditional ship-fighting role.

The U.S. has the world's largest fleet of nuclear-powered attack submarines (54) and nuclear-powered ballistic-missile subs (16). Russia is next, with 37 attack submarines and 14 ballistic missile subs. Britain has 15 nuclear-powered submarines, followed by France with 10, and China with six. Not only are U.S. submarines more numerous but they are also more advanced, the most sophisticated being three 1990s-vintage Seawolfs described by one defense analyst as "the fastest, quietest, and most heavily armed undersea vessels ever built."

Because of the growing power of each of its vessels and the lack of competitors, the U.S. Navy has consolidated its high seas hegemony even while its fleet has shrunk from almost five hundred ships in the 1980s to fewer than three hundred in the early years of the twenty-first century. The potency of U.S. naval vessels is increased by linking together sensors and weapons systems with a tactical Internet known as FORCEnet.

While the U.S. Navy probably will remain unchallenged in blue waters, it faces greater threats as it gets closer to shore. Here water currents, thermal layers, and various obstacles can interfere with even the most advanced sensors, and a variety of defensive weapons systems lurk in wait.

More than seventy-five thousand antiship missiles are owned by seventy countries. A few are ballistic, but most are of the cruise-missile variety. Their potency was proved in 1987 when French-made Exocets fired by an Iraqi aircraft crippled the USS *Stark*, a *Perry-class* frigate, killing thirty-seven sailors. Earlier, Argentinian-owned Exocets sank two British ships during the Falkland Islands War (1982). Newer antiship cruise missiles such as the Russian-made Yakhont, Sunburn, and Uran are even deadlier because they have faster speeds, greater stealth capabilities, and more accurate, GPS-enhanced targeting. Russia is selling these missiles to customers abroad and some nations, such as China, are developing their own versions.

U.S. warships have sophisticated defensive systems to guard against air attack: Incoming missiles can be deflected by electronic countermeasures, flares, or chaff, or destroyed by naval aircraft, sea-to-air Standard missiles, or, as a last resort, by rapid-fire, radar-guided Phalanx guns. But, like the *Stark*, a warship could be caught by surprise or overwhelmed by a flurry of missiles coming from different directions.

Even more worrisome from an American viewpoint is the fact that transport ships and fuel tankers which have to replenish a fleet at sea have no protection when they are outside the defensive range of a battle group. They are as vulnerable as supply convoys on the roads of Iraq. Because a supercarrier has only about a three-day stockpile of JP-5 jet fuel (6,500 barrels a day are needed during combat operations), the most powerful warship in history could be rendered useless if its fuel tankers were sunk. The cargo container ships upon which global commerce depends also have no antimissile protection. A country or even a subnational group armed with cruise missiles could wreak havoc on the world's shipping lanes at chokepoints like the Straits of Hormuz or Malacca.

The threat to shipping, civil and military, is increased by diesel submarines. The latest diesel submarines have ultraquiet electric engines that make them hard to detect with sonar, and they

are much cheaper to buy or produce than a nuclear-powered submarine. Russia has exported *Kilo-class* diesel-electric subs to China, India, Iran, and Algeria, among others. China is producing its own *Song-class* diesel submarines in a bid to challenge U.S. naval hegemony by using the same strategy that Germany, with its U-boats, once used to challenge British dominion of the waves. U.S. antisubmarine defenses are quite sophisticated, especially in open waters, but even American sensors can have trouble tracking quiet diesel subs in noisy coastal waters.

Mines, which can be scattered by submarines or other vessels, represent another major threat to shipping. More than three hundred different varieties are available on the world market. They can be triggered by changes in magnetic fields, acoustic levels, seismic pressure, or other factors. Some come equipped with microelectronics that allow them to distinguish between different types of ships, while others have small motors that allow them to move around. This makes it difficult to certify that a shipping channel is free of mines—it may have been safe an hour ago, but not anymore. Demining technology has lagged behind; the U.S. Navy, for one, has never placed much emphasis on lowly minesweepers. It has paid a price for this neglect. In 1987, during operations to prevent Iran from closing the Persian Gulf, an Iranian mine of World War I design nearly sank the frigate *Samuel Roberts*. Four years later, in the Gulf War, the cruiser USS *Princeton* and the amphibious landing ship *Tripoli* were nearly blasted apart by Iraqi mines. And even a cheap motorboat packed with explosives can pose a significant threat to a modern warship. The USS *Cole*, an *Arleigh Burke-class* destroyer, was badly damaged in such a terrorist attack in the port of Aden, Yemen, in 2000.

All of these threats could be largely negated if U.S. fleets were to stay far out at sea, but they have to approach fairly close to land to launch aircraft or missiles with operational ranges of only a few hundred miles. Moreover, the places where the U.S. Navy is likely to fight in the future are dangerously narrow. The Persian Gulf is only thirty miles wide at its narrowest point, the Taiwan Strait only one hundred miles wide.

To maintain its dominance, the U.S. Navy regularly updates the electronics and weapons aboard its warships even as the hulls and propulsion systems remain unchanged. It also plans to build a variety of unmanned vessels along with a CVN-21 aircraft carrier to replace the *Nimitz-class*, a *Zumwalt-class* DD(X) destroyer to replace *Oliver Hazard Perry-class* frigates and *Spruance-class* destroyers, a CG(X) cruiser to replace the *Ticonderoga-class* cruisers, and a smaller and speedier Littoral Combat Ship with no direct parallel in today's fleet that would focus on clearing mines, hunting submarines, and fighting terrorists in coastal waters. All of these new vessels will have improved defenses and information-processing tools as well as "plug and play" capacity that will allow them to be quickly reconfigured for different missions. They will also incorporate composite materials, stealthier designs, and electric propulsion to make them harder to detect, though an aircraft carrier with a 4.5-acre flight deck can never exactly hide.

Whether all of these warships are truly needed, given the U.S. Navy's already substantial lead over all competitors, remains an open question. A program to develop giant sea bases—perhaps akin to offshore oil-platforms—that would allow American ground and air forces to operate overseas might be of greater use, given the growing difficulty the U.S. has had in gaining basing and overflight rights from other countries.

Aerial Warfare

Fighters such as the American F-15 and the Russian MiG-29 were designed in the 1970s for air-to-air combat, but this has become almost as rare as ship-to-ship actions. Since the Israelis destroyed much of the Syrian air force in 1982, and the U.S. and its allies made similarly quick work of the Iraqi air force in 1991, few if any aircraft have been willing to challenge top-of-the-line Western militaries. (The U.S. Air Force has not produced an ace—an airman with at least five aerial kills—since 1972.) That may change with the sale to China of the Russian-built Sukhoi Su-30, whose performance characteristics are said to exceed those of the F-15C, but the F/A-22 Raptor, the F-35 Joint Strike Fighter,

and the Eurofighter should restore the Western edge. The odds of future aerial dogfights are, however, slim.

Modern surface-to-air missiles pose a more immediate danger, because they are cheaper and easier to operate. The U.S. and its allies have developed effective methods of neutralizing most existing air defenses. In addition to jammers, radar-seeking missiles, and decoys, the U.S. employs stealth technology, first used on the F-117 Nighthawk, then on the B-2 Spirit, and now on the F/A-22 and F-35. Future aircraft may be designed with "visual stealth" technology to render them almost invisible even in daylight.

No other nation has deployed any stealth aircraft. But advanced sensor networks may be able to detect first-generation stealth planes. The Serbs actually managed to shoot down an F-117 in 1999.

None of the most sophisticated surface-to-air missiles, such as Russia's double-digit SAMs (SA-10, SA-15, SA-20), was available to Iraq, Serbia, Afghanistan, or other states that the U.S. has fought in recent years, but they are being sold to other customers. So are shoulder-fired antiaircraft missiles such as the American FIM-92 Stinger, British Starstreak, French Mistral, Chinese Qianwei-2, and Russian SA-7 Grail, SA-14 Gremlin, SA-16 Gimlet, and SA-18 Grouse. There are at least one hundred thousand such systems in the arsenals of over one hundred states and at least thirteen nonstate groups such as Hezbollah, the Revolutionary Armed Forces of Colombia (FARC), and the Tamil Tigers. The best models have a range of 23,000 feet.

The potential of hand-carried missiles was demonstrated in the 1980s when Stingers took a significant toll on Soviet aircraft in Afghanistan. The threat is serious enough for the U.S. to rely increasingly on unmanned drones for high-risk missions and to mandate that manned aircraft in war zones stay above 15,000 or 20,000 feet. SAMs pose an especially great threat to helicopters, which don't have the option of flying that high, and for airplanes taking off or landing. Three cargo aircraft leaving Baghdad International Airport have been seriously damaged by missiles, and, while all of them survived, several U.S. helicopters hit with SAMs in Iraq and Afghanistan did not. An Israeli jetliner was almost shot down in Mombasa, Kenya, in 2002 by al Qaeda operatives firing an SA-7. Only the terrorists' targeting error prevented the deaths of 271 passengers and crew. Other civilian airliners are sure to be less lucky.

Assuming that warplanes can reach their destination, the growing precision of bombs and missiles has made it possible to take out targets with fewer and smaller munitions than ever before. (The U.S. Air Force's latest bomb carries only fifty pounds of explosives.) Weapons are getting smarter all the time. The U.S. Sensor-Fuzed Weapon, first employed in the Iraq War, disperses forty "skeet" antiarmor warheads that use infrared and laser sensors to find and destroy armored vehicles within a thirty-acre area. The Tactical Tomahawk, which entered production in 2004, can loiter up to three hours while searching for targets and receive in-flight retargeting instructions.

The U.S. preponderance in smart bombs and missiles helps to compensate for the relatively small size of its manned bomber force. As of 2005, the U.S. Air Force had only 157 long-range bombers (B-52s, B-1s, B-2s), a considerable reduction not only from World War II (when the U.S. had 34,780) but also from the end of the Cold War (360). While few in number, each B-2 can perform the work of thousands of B-29s by "servicing" eighty "aim points" per sortie.

Tankers such as the KC-10 and KC-135 vastly extend the range and effectiveness of combat aircraft. Cargo-lifters like the U.S. C-5, C-17, and C-130 and the Russian An-70 and An-225 also perform a valuable, if unglamorous, role in projecting military power around the world. The U.S. owns 740 tanker aircraft and 1,200 cargo aircraft, far more than any other country. A lack of such support aircraft makes it difficult for even the relatively sophisticated European militaries to move their forces very far.

A host of other aircraft, ranging from JSTARS and AWACS to Rivet Joint and Global Hawk, perform surveillance and electronic-warfare missions in support of combat forces. Their numbers have been growing: While there were only two JSTARS

in the Gulf War, in the Iraq War there were fifteen. But commanders have become so dependent on these systems that there never seem to be enough to go around. These, too, are vital U.S. assets that few other nations possess.

Space Warfare

A growing amount of surveillance, communications, and intelligence work is being performed by unmanned aircraft and satellites. In 2001 the U.S. had an estimated 100 military satellites and 150 commercial satellites in orbit, as much as the rest of the world combined. The U.S. spends more than $15 billion a year on space, some 90 percent of the global total. The most advanced U.S. surveillance satellites are said to be able to pick out a six-inch object from 150 miles above. A new generation of satellites utilizes stealth technology to make it more difficult for other countries to track their progress and thus to hide equipment from American eyes.

The advantage the U.S. military derives from mastery of space is slowly eroding, however. The Global Positioning System developed by the Defense Department is now widely available for countless commercial applications that have spawned a $30-billion-per-year industry. A potential enemy could use GPS signals to locate targets in the U.S. the same way the U.S. military uses it to locate targets in Iraq or Afghanistan. The U.S. could jam or degrade GPS signals in wartime, but it would have to do so very selectively for fear of imposing a severe toll on the economy, because GPS devices are now essential for civil aviation, shipping, and other functions. In any case, the European Union in cooperation with China is launching its own GPS constellation, known as Galileo, that would be outside of U.S. control.

This is part of a trend of more and more countries—at least forty to date—lofting their own satellites. But getting access to space no longer requires having your own satellite. A growing number of private firms with such names as Google Earth, Keyhole, DigitalGlobe, and Space Imaging sell or give away high-resolution satellite photos via the Internet. The best of these offer imagery of sufficient quality to identify objects eighteen inches wide. The Israeli-owned ImageSat International offers customers the opportunity to redirect its EROS-A imaging satellite (launched in 2000 aboard a Russian rocket) and download its data in total secrecy with few if any restrictions. Its CEO boasts, "Our customers, in effect, acquire their own reconnaissance satellite … at a fraction of the cost that it would take to build their own."

The private satellite industry is becoming so pervasive that the U.S. military has turned to it to provide some of its own imaging (typically low-resolution pictures used for mapping) and much of its communications needs.

Targets identified from space could be attacked either with terrorist (or commando) strikes or with the growing number of missiles proliferating around the world. More than two dozen nations have ballistic missiles and by 2015 at least a dozen will have land-attack cruise missiles. Either type of projectile could be topped with chemical, biological, or nuclear warheads. At least eight countries already have nuclear weapons and more are trying to get them, in part to offset the tremendous U.S. advantage in conventional weaponry.

In response, the U.S. is working on a variety of missile defenses. The most advanced are the ground-based Patriot Advanced Capability 3 and the sea-based Standard Missile 3, which already have been deployed to protect U.S. troops overseas. The deployment of a long-heralded system designed to protect the U.S. homeland against long-range missiles began in 2004 with the installation of interceptors in Alaska. Eventually, the U.S. plans to field a multilayered defense using a variety of sensors and weapons on land, sea, air, and space. Also in the works are systems designed to defeat low-flying cruise missiles, which are hard to distinguish from ground clutter.

FOG FORECAST

The changes wrought by the Information Revolution are still in their early stages and have serious limitations. Even the best surveillance systems can be thwarted by simple countermeasures like camouflage, smoke, and decoys, by bad weather, or by difficult terrain like the deep sea, mountains, or jungles. Sensors have limited ability to

penetrate solid objects, so they cannot tell what is happening in underground bunkers such as those that North Korea and Iran use to hide their nuclear weapons programs. Urban areas present a particularly difficult challenge: There are far more things to track (individuals) and far more obstructions (buildings, vehicles, trees, signs) than at sea or in the sky—and figuring out whether a person is a civilian or an insurgent is a lot harder than figuring out whether an unidentified aircraft is a civilian airliner or an enemy fighter. It is harder still to figure out how hard enemy soldiers will resist or what stratagems they will employ. No machine has yet been invented that can penetrate human thought processes. Even with the best equipment in the world, U.S. forces frequently have been surprised by their adversaries.

Some strategists expect that advances in information technology will greatly diminish if not altogether obliterate some of these difficulties. The Pentagon is creating a Global Information Grid that will pool data from all U.S. assets, whether an infantryman on the ground or a satellite in space. The ultimate goal: to provide a perfect operational picture—a "God's-eye view" of the "battlespace."

This ambitious objective could be furthered by the development of better microwave radars that could see through walls, foliage, or soil; cheaper, more pervasive sensors that could provide 24/7 coverage of the battlefield; better data compression and transmission techniques that could allow more bytes to be sent much faster; and more powerful computers that might make it possible to create, for example, a real-time, three-dimensional model of a city showing all the people who reside in it. Not only are electronic computers getting more powerful in accordance with Moore's Law, but even greater leaps in computing power may be achieved using biological computers that rely not on electronic components but on DNA, individual molecules, or proteins; optical computers that use light rather than electricity; and quantum computers that exploit the atomic properties of matter. Some scientists speculate that by using such methods a billionfold increase in computing power is possible. "That means you could hold the power of all earth's present computers in the palm of your hand," says Stan Williams of Hewlett-Packard.

Yet no matter how far information technology advances, it is doubtful that the Pentagon will ever succeed, as some Utopians dream, in "lifting the fog of war." The fallibility of American soldiers and the cunning of their enemies will surely continue to frustrate their best-laid plans. Indeed, America's growing reliance on high-tech systems creates vulnerabilities of its own: Future enemies have strong incentives to attack U.S. computer and communication nodes. Strikes on military information networks could blind or paralyze the armed forces, while strikes on civilian infrastructure, such as banking or air-control systems, could cause chaos on the home front.

Safe to say, adversaries will figure out ways to blunt the U.S. informational advantage. From Operation Anaconda in Afghanistan to numerous misadventures in Iraq, they already have.

AMERICA'S UNPARALLELED POWER

While various setbacks in the war on terror underscore the limits of American power, it is important not to lose sight of the bigger picture: The dawn of the Information Age has coincided with a period of American global hegemony. Part of the explanation lies in U.S. economic strength and its "soft power." But Europe and Japan are rich, too, yet their global sway lags far behind. What they lack is America's military strength.

In the early years of the twenty-first century the United States enjoys a preponderance of military power greater than that of any other nation in history. Rome was rivaled in land power by Carthage, Persia, and other states, and its sphere of control never extended far beyond western Europe and the Mediterranean region. Britain was rivaled in naval power by the Netherlands, France, and later Germany, and it was never dominant on land. Today America is rivaled in land, sea, and air power by … no one. Although the dominance of U.S. forces can still be challenged when they come into close contact

with the enemy on his home turf, they are undisputed masters of the "commons" (sea, air, space), which allows them to project power anywhere in the world at short notice. In the words of journalist Gregg Easterbrook, "The American military is now the strongest the world has ever known, both in absolute terms and relative to other nations; stronger than the Wehrmacht in 1940, stronger than the legions at the height of Roman power."

What accounts for this huge military edge? Not manpower. The U.S. armed forces in 2006, with about 1.4 million active-duty personnel and 1.2 million reservists, were hardly the biggest in the world. NATO's European members and China each had roughly 2.2 million men and women under arms. Other nations—North Korea with 1 million active-duty soldiers and 4.7 million reservists, Russia with 960,000 on active duty and 2.4 million in the reserves—were not far behind. In any case, because the U.S. armed forces typically do their fighting overseas, they seldom enjoy a numerical advantage in their theater of operations. Whether in Iraq or Afghanistan, they have usually been outnumbered.

In the Information Age the U.S. edge lies in quality, not (as it did during the Second Industrial Age) in quantity. Both its soldiers and their equipment are among the best on earth. The U.S. armed forces are not the only Information Age military—Britain, France, Japan, Singapore, Australia, Sweden, and other nations have followed in its wake—but they have gone the furthest fastest.

Among other nations, Israel's success in harnessing the Information Revolution has been particularly consequential. The balance of power in the Middle East shifted sharply in Israel's favor following the Yom Kippur War, because it modernized its military to a far greater extent than its neighbors did. This helps to explain why no Arab nation has mounted a conventional attack on Israel since 1973.

The United States' success in harnessing the Information Revolution produced its profoundest effects without a shot being fired. Aware that the Soviet Union could not keep up with U.S. military advances in the 1980s, Mikhail Gorbachev responded with perestroika and glasnost, reform efforts that spun out of control and ultimately consigned the Soviet Union to the ash heap of history. The demise of the USSR set the stage for the U.S. military dominance that has underpinned the current era of globalization, which has been remarkably devoid of conflict between states.

Wars between states have declined in size and frequency for a variety of reasons, not the least being the integration of old adversaries like France and Germany into the European Union, the international norms propagated by the United Nations Charter, the spread of capitalism and democracy across the world, and the deterrence provided by nuclear weapons. But surely another part of the explanation is that few states feel free to challenge the Pax Americana. Several of those that have tried, including Panama, Serbia, Afghanistan, and Iraq, have seen the downfall of their regimes. "War, let alone 'decisive war' between major states, currently is enjoying an off season for one main reason," writes British strategist Colin Gray. "So extreme is the imbalance of military power in favor of the United States that potential rivals rule out policies that might lead to hostilities with the superpower."

ACCESS DENIED?

It is doubtful that any country will mount a full-spectrum challenge to U.S. military capabilities in the foreseeable future. The entry barriers are simply too high, especially for air, sea, and space systems. *Virginia-class* nuclear submarines cost $2.4 billion, *Nimitz-class* aircraft carriers go for $6 billion, and the F-35 Joint Strike Fighter program will cost at least $245 billion. The U.S. spends around $500 billion a year on its military, almost as much as the rest of the world combined. In fact, the U.S. spends more simply on the research, development, testing, and evaluation of new weapons—$71 billion in 2006—than any other country spends on its entire armed forces. (By way of comparison, the top three spenders after the U.S. are Russia, whose defense budget in 2003 was estimated at $65 billion; China, at $56 billion; France, at $45 billion; and Japan and the United Kingdom, at $42 billion each.)

It is not only U.S. hardware that's hard to replicate; so is the all-volunteer force that makes it work. Operating high-tech military equipment requires long-service professionals, not short-term conscripts. Countries as diverse as Vietnam, China, Germany, and Russia are emulating the Anglo-American model by downsizing their forces and relying less on draftees. Many other nations have abolished the draft altogether. The U.S. military's edge lies not simply in recruiting high-quality personnel but in its methods for training and organizing them. Initiatives undertaken in earlier decades, such as setting up realistic training centers to simulate combat conditions and forcing the services to work more closely together (the Goldwater-Nichols Act), continue to bear fruit. Few other armed forces have made comparable reforms. Writes one scholar: "The superb U.S. all-volunteer military force … is a unique human and institutional asset that less capable foreign rivals can neither copy nor steal."

But a potential adversary does not need to duplicate U.S. force structure in order to challenge it. The United States faces a growing "asymmetric" threat both from other states and from substate groups.

As we have seen, a variety of off-the-shelf missiles can threaten U.S. tanks, surface ships, and aircraft, especially when they get close to hostile territory. The power of smart munitions is outstripping the protection afforded by speed or armor. After 2010, write defense analysts Michael Vickers and Robert Martinage, "the survivability of aircraft carriers, high-structure surface combatants [e.g., tanks], and non-stealthy aircraft of all types could increasingly be called into question as maritime, over-the-horizon 'area denial' capabilities and extended-range air defense systems continue to mature."

Also vulnerable are the ports, airfields, and bases that the U.S. uses to project its power overseas. Imagine how much damage Saddam Hussein could have done in 2003 if he had been able to annihilate the one port in Kuwait that was being used to disembark coalition troops or the large desert bases in Kuwait where over one hundred thousand British and American troops gathered prior to the invasion of Iraq. The Pentagon's 2001 Quadrennial Defense Review warned that "[f]uture adversaries could have the means to render ineffective much of our current ability to project military power overseas."

If the U.S. armed forces could not count on safe, assured access to overseas bases they would have to radically change the way they do business. It would no longer be practical to rely on large land armies or lots of short-range combat aircraft operating out of vulnerable forward bases supplied by equally vulnerable cargo ships, trucks, and aircraft. The U.S. Army might be forced to rely on a small number of commandos supported by long-range aircraft and missiles—the Afghanistan model. The navy might have to depend more on submarines and the air force on stealth aircraft. All the services might have to make greater use of unmanned vehicles. The battlefield, which has been growing less crowded for centuries, might empty out even further as small units try to conceal themselves from ubiquitous sensor networks, emerging only briefly to launch lightning strikes before they go back into hiding.

This has become known as the "swarming" scenario, and it has attracted support from the likes of military historian Alexander Bevin, who writes, "Large concentrations of troops and weapons are targets for destruction, not marks of power, and [in the future] they no longer will exist.… Military units, to survive, must not only be small, but highly mobile, self-contained, and autonomous."

Even if these predictions are accurate, however, no timeline is attached. Timing matters. After all, everyone knows that a bull market will be followed by a bear market, but the key to making money in stocks is knowing when that transition will occur. Likewise, the key to winning future wars is knowing when to move from one form of military to another. A premature decision to change (such as the U.S. Army's Pentomic design in the 1950s) can leave one unprepared to fight and win wars that actually occur, Vietnam being the classic example.

In any case, it is doubtful that a complete shift to "swarming" will ever take place. Winning wars, as distinct from winning battles, will continue to require controlling territory, which in turn will

require a substantial presence of ground troops, as the U.S. has learned in Afghanistan and Iraq. No wonder weapon will alter this fundamental reality.

"AMERICAN HIROSHIMA"

Even as strategists look to the future, armed forces must not lose sight of the threats of the moment, and they do not come for the most part from traditional militaries. They come largely from terrorist groups—some with state sponsorship, others without—that utilize the fruits of modern technology to their advantage.

"Irregular" attacks carried out by tribes, clans, or other nonstate actors are as old as warfare itself; they long predate the development of modern armed forces and the nation-state. The religious fanaticism that animates so many of today's terrorists and guerrillas is equally ancient. But technological advances have made such attacks far more potent than in the past. The progeny of the Second Industrial Revolution—assault rifles, machine guns, mortars, rocket launchers, land mines, explosives—long ago spread to the remotest corners of the globe. Fighters who a century ago might have made do with swords and muskets now have access to cheap and reliable weapons, such as the AK-47, capable of spewing out 100 bullets a minute. More advanced technologies, from handheld missiles to chemical, biological, and nuclear weapons, give even a small group of insurgents the ability to mete out far more destruction than entire armies could a century ago. And thanks to modern transportation and communications infrastructure—such as jumbo jets, the Internet, and cell phones—insurgents have the capability to carry out their attacks virtually anywhere in the world.

September 11, 2001, showed the terrifying possibilities of such unconventional warfare. It is easy to imagine that in the future super-terrorists will be able to kill hundreds of thousands, even millions, with effective weapons of mass destruction. All of the materials, as well as the know-how, needed to craft such devices are all too readily available…

The atomic bomb is more than sixty years old. It belongs to an age of rotary-dial telephones and fin-winged cars. It is a miracle that it has not been used in anger since 1945, but that streak won't last forever. And while Information Age technology offers a reasonable chance of stopping a nuclear-tripped missile, there is much less probability of stopping a terrorist with a nuclear suitcase. (Yes, nuclear bombs have gotten smaller, just like computers.) There is little in theory to prevent al Qaeda from carrying out its oft-expressed desire to create an "American Hiroshima." In the words of a retired four-star general who once ran antinuclear terror programs for the Department of Energy: "It is not a matter of *if*, it's a matter of *When*."

The most important challenge for the armed forces of the U.S. and its allies in the post-9/11 world is to "leverage" their advantage in conventional weaponry to deal with an unconventional threat. Information technology can be an important part of this task. Embedded microchips can track the 18 million cargo containers moving around the world and prevent terrorists from using them to smuggle weapons. Computerized cameras scanning a crowd may be able to pick out a terrorist based on facial recognition patterns. Doglike sniffing machines may be able to recognize suspects by their body odor. Powerful computers utilizing artificial intelligence programs can sift vast reams of digital data to pick out information about terrorist plots—if concerns about violating the privacy of innocent people do not get in the way. A variety of unobtrusive sensors can warn of the presence of explosives or chemical, biological, or nuclear weapons. Handheld computer translating devices such as the Phraselator already in use by U.S. troops can bridge some of the language gap between Western operatives and the regions where they operate.

But in the final analysis, having the best technology is not enough to defeat the most ruthless terrorists. In fact, most of the expensive weapons systems being purchased by the U.S. and its allies are almost completely irrelevant to the war against terrorism. Smart bombs can be useful for killing the enemy once he's been located. But figuring out who the enemy is, where he is, and what he is up to—that requires smart *people*.

Various remedies have been proposed to fix the obvious American weakness in this area: Requiring more foreign-language training for the armed services. Recruiting more spies. Setting up a military school devoted to civil governance. Creating an agency tasked with "nation-building." Appointing "cultural scouts"—an enhancement of the existing Foreign Area Officer program, which is widely seen as a career dead end—who would spend long periods of time in foreign lands. These are all good ideas and some are beginning to be implemented, but they are harder to enact than a decision to buy a new weapons system. They require changing the culture of the Defense Department and other government agencies, which have not stressed skills in counter-insurgency, human intelligence, or nation-building.

TOOTH AND TAIL

While the U.S. armed forces have made some major strides in recent decades, especially in getting branches to work more closely together, they still have a long way to go before they have an organizational structure that makes the most effective use of their high-tech equipment. The army, although putting greater emphasis on more nimble, four-thousand-person brigades, still maintains a divisional structure that dates back to the First Industrial Age, while the basic units of the navy (carrier groups) and the air force (squadrons and wings) date back to the Second Industrial Age. All of the services employ a personnel system developed more than one hundred years ago that treats individuals as interchangeable cogs in a giant machine. Army officers have traditionally spent an average of only eighteen months at each duty station during the course of a twenty-five-year career, making it difficult to build true expertise in such complex areas as foreign cultures.

Another problem is a cumbersome procurement process that costs too much and takes too long to deliver new equipment. The U.S. paid a heavy human cost for those shortcomings in Iraq, where body armor, armored Humvees, and IED-jamming devices were slow to arrive. The time and expense involved in producing new weapons systems have increased severalfold since the 1970s, even as the private sector has been able to produce products like cell phones and microchips faster and more cheaply. "In order to stay ahead of adversaries with access to technologies available in the global marketplace, the DOD [Department of Defense] needs to shorten the timeframe from concept to testing," writes Lieutenant Colonel Augustus Way Fountain III, a chemistry professor at West Point. That's exactly what the best-run companies—the Microsofts and the Wal-Marts—have done, but the U.S. armed forces (and most others, for that matter) have lagged behind.

In general, the U.S. military is hindered by a sluggish, bloated bureaucracy that has resisted countless reform efforts. "We're No. 1 in the world in military capabilities," says David, M. Walker, head of Congress's Government Accountability Office. "But on the business side, the Defense Department gets a D—giving them the benefit of the doubt. If they were a business, they wouldn't be in business."

Some statistics buttress the point. The "tooth to tail" ratio in the U.S. military—the number of support personnel as opposed to actual trigger-pullers—varies from 12:1 for the Army and Marine Corps, to 15:1 for the Navy, and 32:1 for the Air Force. "An astounding 70 percent of the defense budget is spent on overhead and infrastructure (the bureaucratic 'tail'). Only 30 percent directly reaches our combat forces in the field (the 'tooth')," writes retired Admiral Bill Owens. "No community would tolerate 7 out of every 10 police officers sitting at their desks pushing paper. The nation should not tolerate such a ratio in the military."

What makes this unwieldy setup so intolerable is that it can stymie the potential of even the most advanced information technologies. Retired Marine Colonel T. X. Hammes notes that most commanders must submit their intelligence requirements up the chain of command. Each level validates, consolidates, and prioritizes the requests, which are then fed through the centralized staff system to task the assets that will actually collect against the requests. The information is collected, passed to another section for analysis, then put in the form of a usable product, and finally disseminated through the same cumbersome system. Thus, the premier benefit of

the Information Age—immediate access to current intelligence—is nullified by the way we route it through our vertical bureaucracy.

The good news, from the American perspective, is that many potential enemies (e.g., China or North Korea) are encumbered by even more sclerotic bureaucracies. The bad news is that this is not true of nimble, networked groups like al Qaeda. To fight them effectively, the U.S. military will have to display more of the decentralized decision-making that it showed in Afghanistan in the fall of 2001. This will not be easy to do because modern command and control technology is a two-edged sword: It can be used to centralize or decentralize. It will always be difficult for a senior military or political leader to resist the temptation to micromanage operations from afar—a style of leadership that modern communications technology has made easier, but no more effective, than in the past.

Veteran journalist Robert Kaplan rightly calls "the dinosauric, vertical bureaucracy of the Industrial Age … the greatest single impediment to America's ability to wage a successful worldwide counter-insurgency." Unless the U.S. government can streamline its Industrial Age bureaucracy and become a networked organization, it may find that even purchasing the latest and best technology will not offer sufficient protection against the country's foes.

2.1.2 THE AUTOMATION AND PROLIFERATION OF MILITARY DRONES AND THE PROTECTION OF CIVILIANS

NOEL SHARKEY

Most people associate unmanned aerial vehicles (UAVs) or drones with military surveillance and intelligence gathering. However, increasingly these vehicles are becoming weaponized, and are being used for combat. On September 30, 2011, an American citizen, Anwar al-Awlaki, who had called for jihad against the United States, was killed by an armed Predator drone in Yemen. Two weeks later his sixteen-year-old son, who was also an American citizen, was also killed in a CIA drone strike.

This selection details the development and proliferation of unmanned aerial combat vehicles (UACVs) and raises some troubling legal and ethical issues about the reliance on these kinds of weapons. There is currently an arms race taking place in the development of UACVs, with America and Israel leading, for the time being. But with other nations working to develop their own UACVs, it is likely that these kinds of weapons will soon be in the arsenals of many more nations.

More troubling are the ethical and legal concerns that the potential use of these weapons raise. Will autonomous drones and robots be capable of conforming to international laws of war that require that targeting strategies distinguish between combatants and noncombatants, and that the force that is employed be proportional to the threat? Are we heading into a future in which human lives can be snuffed out by algorithms?

Source: Source: From Noel Sharkey, "The Automation and Proliferation of Military Drones and the Protection of Civilians," *Law, Innovation and Technology*, December 2011, Volume 3, Number 2, pp. 229–240. Reprinted by permission of Hart Publishing.

FOCUS QUESTIONS

1. The author raises several important legal questions concerning the use of remotely controlled armed aircraft, in particular, whether the introduction of these vehicles into a war zone constitutes the introduction of armed forces or an "act of war." What is your view on this question? Defend your answer.
2. The author expresses concern that the military use of autonomous drones could violate the laws of war. What reasons does he give for thinking this is a problem? Do you agree? Explain.
3. Under what kinds of circumstances have the indiscriminate killings of civilians been justified in times of war? Is it conceivable that similar circumstances could be used to justify the deployment of unmanned autonomous drones or robots? Explain.

KEYWORDS

armed forces, autonomous drones, indiscriminate weapons, international humanitarian law, Karrar, military necessity, moral superiority, non-combatants, Predator, proliferation, Reaper, Switchblade, targeted killings, UAV, UCAV

I. INTRODUCTION

In the last decade, the growing use of the Unmanned Aerial Vehicle (UAV) for intelligence, surveillance and reconnaissance has been a game changer on the battlefields of Iraq, Afghanistan and the Middle East. The US is leading the technology, with Israel in second place. There were only 150 robots in the Iraq War in 2004, including bomb disposal robots. Now there are thousands of small, unarmed aerial surveillance drones being used by troops on the ground—so many, in fact, that it is difficult to obtain an accurate estimate of their number.

The main concern of this paper is the use of drones armed with missiles and bombs. Missions in Afghanistan and Iraq are mainly flown by 'pilots' of the 432nd Air Expeditionary Wing at the Creech Air Force Base in the Nevada desert, thousands of miles away from the operations. Operators sit at video consuls making decisions about when to apply lethal force. Because these (unmanned) planes can be flown around the clock, the pilots can take a break from 'battle' at any time or even go home for dinner with their children. This unusual combination of spending all day in the theater of war and then returning home in the evening is apparently causing a new kind of battle stress that has not been known before. The argument is that military personnel at war normally spend their time with their comrades and this reduces their stress (but the jury is still out on this until substantial evidence accrues).

MQ-1 Predators, unmanned Combat Air Vehicles (UCAV) which each carry a payload of two hellfire missiles, clocked up 750,000 miles in the air between 2007 and 2010. In October 2007, the Predator was joined by the much larger and more powerful MQ-9 Reaper. The MQ-9 Reaper carries a payload of up to 14 hellfire missiles or a mixture of missiles and bombs. These so-called hunter-killer UAVs have conducted many decapitation strikes since they were first deployed in Afghanistan in October 2007.

Despite cutbacks in US defense funding, the Obama administration has continued to increase its spending on UCAVs. In 2010 the Air Force aimed to spend $2.13 billion on unmanned technology, with $489.24 million earmarked to procure 24 new heavily armed Reapers. The US Army planned to spend $2.13 billion on unmanned vehicle technology. This includes the purchase of 36 more unmanned predators. The US Navy and Marine Corps planned to spend $1.05 billion on unmanned vehicles, including armed MQ-8B helicopters.

US drone developments are moving quickly, with plans well underway to arm both the Fire Scout unmanned helicopter and the small rail-launched Shadow RQ7-B (approx 4.6 m long). Aeroenviroment, a US company, has developed the Switchblade, a small 2 lb backpack UAV launched from a tube. It unfolds its wings and streams video back to the controller. When a target is detected, the controller designates it and the Switchblade folds its wings back to become an autonomous missile.

Israel is the second largest producer of UAVs and is continually enhancing its military strike capabilities and range. China's plans for new concept UAVs are also coming along very strongly. Many other countries are following suit. The first step for most countries will be to arm their UAVs, and the next step will be to automate them. Again, the US and Israel are the leaders.

Before we examine the ethical and legal implications of the new robot weapons, we will first take a broad sweep over some of the recent history of UAV development and proliferations. We will then look at the current use of UAVs before discussing the near-future of autonomous supersonic (and faster) air craft and how these developments could impact on civilian populations if they remain unchecked. Finally, we will express concerns about how military necessity could unlock the use of autonomous war planes even though they remain indiscriminate weapons.

II. PROLIFERATION

In the post-9/11 era, aerial drones have come to dominate military operations. Troop movements are almost always accompanied by intelligence, surveillance and reconnaissance drones. The military success of UAVs in the theater of war in Iraq and Afghanistan create massive worldwide demand for the technology. It is massive business. The Teal Group has estimated that the market will increase to $11.3 billion per year over the next decade, not including the billions of dollars in development costs. Teal does not have access to the figures for military robotics spending from major countries such as Russia, China or Iran.[1]

There are now at least 50 countries using UAVs.[2] Many of these are being developed in-house and many are being bought in (and probably copied). The US sells many of its drones to its closest allies in Europe, and recently the US company General Atomics was granted permission to sell its earlier generation predators in the Middle East and Latin America. Israel has an even wider range of markets, having recently expanded into Latin American countries. Countries that do not have the advantage of being a close ally of the US cannot yet buy armed drones, and so they are having to find other means of acquiring or developing them. India and Pakistan are working hard to develop attack drones, having failed to purchase any from the US or Israel. Russia has shown models of the MiG Skat unmanned combat aircraft, which is intended to carry out strike missions on air defenses. It is, according to reports from Russia, able to carry cruise missiles and can strike both ground and naval targets. Iran demonstrated a rocket-launched UAV, the Karrar or ambassador of death, to the press in 2010. It carries two cruise missiles. It is not possible to ascertain how operational the Iranian and Russian craft are, but it is clear that, at the very least, they are moving in the right direction to make the technology.

China is showing the greatest commercial potential for selling armed UAVs over the coming decade. It has made a showing of many different types of UAV at its air shows over the last five years, some almost replicas of the US drones. The US-China Economic and Security Review Commission said that China 'has deployed several types of unmanned aerial vehicles for both reconnaissance and combat'.[3] According to the *Washington Post*, at the Zhuhai air show in China in November 2010, there were more than two dozen Chinese UAV models on display.[4] Worryingly, the *Washington Post* quotes Zhang Qiaoliang of the Chengdu Aircraft Design and Research Institute as saying, 'The United States doesn't export many attack drones, so we're taking advantage of that hole in the market.'

This is worrying because it indicates the opening up of a large and expanding market of which all the major players will want a share. If it looks like China's combat UAVs threaten to dominate the market, then others will start selling them and every developed nation will have them. This could have a significant impact on how disputes are handled and what constitutes a war.

III. KINETIC ATTACKS AND ACTION SHORT OF WAR

Legal and political loopholes for drones are already being created in US that give us some insight into what the future might hold when every major power is regularly deploying armed robots. Will drones change the definition of what constitutes a war? Will their use be considered to be an act of war, a policing action or action short of war (under *jus ad bellum*)? US operations in Libya have sparked an argument between President Obama and the US Congress over the War Powers Resolution. The US 1973 War Powers Resolution limits the ability of a president to wage war without Congressional approval. The president is required to obtain congressional approval or terminate a mission within 60 days

> [i]n the absence of a declaration of war, in any case in which United States Armed Forces are introduced—(1) into hostilities or into situations where imminent involvement in hostilities is clearly indicated by the circumstances; (2) into the territory, airspace or waters of a foreign nation, while equipped for combat, except for deployments which relate solely to supply, replacement, repair, or training of such forces ...[5]

The President did not seek Congressional approval for the United States' role in NATO's Libya mission within the statutory 60 days. Harold Koh, the most senior lawyer in the state department, has strongly defended the legality of US military involvement in Libya without Congressional approval. The argument was that since April 2011 the USA's role in the NATO-led mission in Libya has mainly involved assisting in operations (refuelling and providing intelligence) and this complies with (2) above, or carrying out missile strikes with armed drones.

A report to the US lawmakers explaining why the President did not need to seek Congressional approval stated: 'US operations do not involve sustained fighting or active exchanges of fire with hostile forces, nor do they involve the presence of US ground troops, US casualties or a serious threat thereof.'[6]

The actual wording of the 1973 War Powers Resolution concerns the introduction of *armed forces* into hostilities or into the territory, airspace or waters of a foreign nation, while equipped for combat. The argument here would have to be that use of drone strikes in Libya did not constitute the introduction of US armed forces into hostilities or into foreign airspace.

There are at least two important questions that need be addressed here. The first is this: are drones now considered to be action short of warfare? As the White House told the *New York Times*: 'American involvement fell short of full-blown hostilities.'[7] The second question is, does the use of remotely piloted armed aircraft constitute the introduction of armed forces or not?

Although the US Congress has only been concerned about involvement in Libya, the use of drones for targeted killing has been carried out in at least three other countries that the US is not at war with: Somalia, Yemen and Pakistan. Other than Libya, US drone operations in these countries are carried out by the intelligence services. Following the withdrawal of the DASH in 1979, the Central Intelligence Agency (CIA) was the first in the US to use armed drones. In 2002 they killed five men travelling in a Sport Utility Vehicle in Yemen.[8]

Department of Defense lawyers considered this to be a legitimate defensive preemptive strike against al-Qaeda by US Department of Defense lawyers. Since then, the use of drones for targeted killings or 'decapitation strikes' in states that are not at war with the US has become commonplace.

Although the US will neither confirm nor deny the strikes officially, the *Asia Times* has called the CIA drone strikes 'the most public "secret" war of modern times'.[9] The former Director of

the CIA, Leon Panetta, has certainly been vocal about the operations. In 2008 he told the Pacific Council on International Policy: 'it's the only game in town in terms of confronting and trying to disrupt the al-Qaeda leadership.'[10] Revealing the CIA's intentions regarding the expansion of targeted drone kills, Panetta went on to say of al-Qaeda that, 'If they're going to go to Somalia, if they're going to go to Yemen, if they're going to go to other countries in the Middle East, we've got to be there and be ready to confront them there as well. We can't let them escape. We can't let them find hiding places.'[11]

Apart from Libya, none of the drone strikes in countries not at war with the US have even been considered by Congress under the War Powers Resolution. This is a dangerous precedent which is, at best, legally questionable under international humanitarian law as pointed out by Philip Alston, UN Special Rapporteur on extrajudicial killings. He challenged the legality of the targeted killings at a UN General Assembly meeting in October 2009. A request was issued for the US to provide legal justification for the CIA's targeting and killing of suspects with a question about who was accountable. The US refused to comment on what they said were covert operations and a matter of national security.

US Department of State legal advisor Harold Koh rebutted Alston indirectly, stating that 'US targeting practices including lethal operations conducted by UAVs comply with all applicable law including the laws of war'.[12] However, there are no independent means of determining how the targeting decisions are being made. It remains unclear as to what type and level of evidence is being used to make decisions that effectively amount to death sentences by Hellfire missiles for non-state actors. The suspects are not provided with an opportunity to appeal or even to surrender. It is also unclear as to what other methods, if any, are exhausted or attempted to bring the suspects to justice. The whole process is taking place behind a convenient cloak of national secrecy.

US law professor Kenneth Anderson questioned the CIA's use of drones in a prepared statement to a US Senate hearing:

> [Koh] nowhere mentions the CIA by name in his defense of drone operations. It is, of course, what is plainly intended when speaking of self-defense separate from armed conflict. One understands the hesitation of senior lawyers to name the CIA's use of drones as lawful when the official position of the US government, despite everything, is still not to confirm or deny the CIA's operations.[13]

A subsequent report by Alston in 2010 to the UN General Assembly[14] describes drone strikes as violating international and human rights law because both require transparency as to the procedures and safeguards that are in place to ensure that killings are lawful and justified: '[A] lack of disclosure gives States a virtual and impermissible license to kill.' Some of Alston's arguments also revolve around the notion of 'the right to self-defense' and whether drone strikes are legal under Article 51.

It appears that the US does not consider the CIA strikes or the deployment of armed drones in Libya as acts of war. How far will this go? All of the countries that have been subject to strikes are militarily inferior and pose little threat to western nations. It seems unlikely that more militarily sophisticated countries such as China or Russia would see such actions on their territory as actions short of war. The precedent is now in place and it will be interesting to see what happens when other countries start doing the same.

IV. AUTONOMY AND THE PACE OF BATTLE

Since 2004, all roadmaps of the US forces have made clear the desire and intention to develop and use autonomous battlefield robots. Execution of these plans to take the human out of the loop is well underway. The end goal is that robots will operate autonomously to locate their own targets and destroy them without human intervention.[15]

Autonomous lethal targeting is not illegal so long as it accords with the principles of distinction and proportionality. In a military rich environment with few civilians in, say, a desert or at sea, there may be few problems with using armed robots to kill targets. Legally this may be little

different from firing rockets from a distance, dropping bombs or sending cruise missiles. However, armed robots are set to change the pace of battle dramatically in the coming decade. It may not be militarily advantageous to keep a human in control of targeting.

The speed of an unmanned craft is limited only by the integrity of its structure and components and not by human G-force limitations. Unmanned planes can not only travel faster than piloted planes but can also maneuver much faster, taking sharp turns that would kill a human pilot.

The US has been testing the supersonic Phantom Ray and the X-47b. The US Navy would like to replace the F-35s on their carriers with the X-47b.[16] The Chinese (Shenyang Aircraft Company) are working on the Anjian (Dark Sword) supersonic unmanned fighter aircraft, the first UCAV designed for aerial dogfights. DARPA and the Pentagon want armed unmanned vehicles that can reach anywhere on the planet within 60 minutes. The DARPA HTV-2 program is a good example of the direction of technological developments. The Falcon HTV-2 is a hypersonic unmanned plane that in recent tests flew at a velocity of between 17 and 22 Mach, ie, 17 to 22 times the speed of sound at its altitude. That is a maximum speed of 13,000 mph (20,921.5 kph), which is around 8.5 times faster than the Russian MiG-25, maximum velocity Mach 2.3 (1,520 mph or 2,446 kph).

However, as with any mobile system controlled by complex software we cannot predict how it will react in all circumstances. A series of unpredictable events could occur, or there could be an undetected bug in the program or a hardware fault. A hypersonic drone could be off target by 5 km in less than a second.

A simple example of just two interacting software algorithms running out of control happened on the Amazon website. The out-of-print 1992 book *Making of a Fly* usually sells for around $50. But on 19 April 2011 Borderbooks was selling it for $23,698,655.93 (plus $3.99 shipping) on the Amazon website.[17] This astonishing price was created because an automatic algorithm from bookseller Profnath was interacting with Borderbooks' automatic algorithm. The story is that when Borderbooks does not have a book in stock, they automatically list it at 1.27059 times the price of the highest other seller. So when a customer orders it from them, they can buy and sell at a profit. The problem was that Profnath's algorithm made their prices 0.9983 times the highest price of other booksellers. So each time Borderbooks increased their price, so did Profnath, and they spiralled out of control.

This was quite harmless, as no one was prepared to pay these kinds of prices. However, imagine two or more complex algorithms interacting on high-speed armed robots. Without any knowledge of the other algorithms, there is no way to tell what would happen. They might just crash into one another or into the ground, or they might end up unleashing their destructive power in completely the wrong place. The point is that software algorithms on autonomous armed drones spiralling out of control is something to be very seriously concerned about.

As I have written elsewhere,[18] allowing robots to make decisions about the use of lethal force could breach both the principle of distinction and the principle of proportionality as specified by international humanitarian law. These are the pillars of the laws of war. Currently and for the foreseeable future no autonomous robots or artificial intelligence systems have the necessary properties to enable discrimination between combatants and civilians or to make proportionality decisions.

Under the principle of distinction, only combatants/warriors are legitimate targets of attack. All others, including children, civilians, service workers and retirees, should be immune from attack. The same immunity covers combatants who are *hors de combat*—those who are wounded, have surrendered or are mentally ill.[19] The principle of proportionality applies in circumstances where it is not possible to fully protect non-combatants in an action. It requires that the loss of life and damage to property incidental to attacks must not be excessive in relation to the concrete and direct military advantage anticipated.

Distinguishing between civilians and combatants is problematic for any robot or computer system. First, there is the problem in the specification

of 'civilian-ness'. A computer can compute any given procedure that can be written as a program. We could, for example, give the computer or a robot an instruction such as, 'if civilian, do not shoot'. This would be fine if and only if there was some way to give the computer a precise specification of what a civilian is. The laws of war don't help. The 1949 Geneva Convention requires the use of common sense to determine the difference between a civilian and combatant, while the 1977 Protocol essentially defines a civilian in the negative sense as someone who is not a combatant.

Two major software components are necessary in order for robots to distinguish between combatants and non-combatants. The first is highly accurate and discriminative sensing and visions systems. While technology has improved dramatically over the past 50 years, the development of software for vision systems has been very slow. Currently we have vision systems that can detect whether something is a human or not by its shape, although these can easily be fooled by a mannequin or a dancing bear. We have face recognition systems that are effective so long as individuals stay still long enough to be identified, and we have various biometric tests for people who are stopped. In the fog of war all of these methods run into insurmountable difficulties.

The second necessary component is reasoning from situational awareness. It is unclear as to when we might even get a foot in the door for this problem. There are always optimists, but the truth is that such systems are in the realm of 'hope ware' rather than software. There is no way to be certain that they will never be achieved, but equally there is currently no evidence to suggest that they will ever be achieved. If they are achieved, it could be in hundreds of years.

In terms of the laws of war, we must go on the information and evidence that we currently have. We should not rely on technological fixes that are just around the very elastic corner that we may never reach. The bottom line is that autonomous robots that can kill without a human making the lethality decisions are indiscriminate weapons. They may properly belong in the United Nations Convention on Certain

V. MILITARY NECESSITY

The statement 'Armed robots will always have a person somewhere in the loop for lethal targeting decisions' is often repeated by Western powers. But saying 'somewhere in the loop' is not the same as saying that a human will always make the lethal targeting decisions. There are clearly instances where military necessity may override having anyone in the loop. In the extreme, if the very survival of the state was at stake and it was possible to use autonomous armed robots to save the day, it is fair to say that they would be used. In these circumstances the use of autonomous killing machines may even be considered a legitimate action—or, at least, such an action might not be considered to be illegitimate if judged relative to the International Court of Justice's decision, or more properly non-decision, regarding the use of nuclear weapons by States. As is well known, the Court ruled that, in the current state of international law and given the facts at its disposal, it was not possible to conclude definitively whether the threat or use of nuclear weapons would be lawful or unlawful in extreme circumstances of self-defense (circumstances in which the very survival of the defending State would be at stake).[21] It would not be too fantastic to imagine the phrase 'autonomous armed robots' being substituted for 'nuclear weapons'. Armed robots are a lesser beast than nuclear weapons—unless they are armed with nuclear weapons of course. So the substitution is easy. However, it is likely that it would take much less than the imminent collapse of a state before indiscriminate autonomous robots were deployed.

History is littered with many examples in which humanitarian considerations have been overridden for the protection of soldiers rather than for the survival of the state from imminent collapse.

The attacks on the French market town of St. Lô during the Normandy invasion by allied forces in 1944 provide a good example of the indiscriminate use of air power. Although the town was full of friendly French civilians, an elite German Panzer division residing there was blocking the

Canadian, US and British forces took very heavy casualties. In response, according to Cohen, 'The town was attacked on 25 July, by 1,500 heavy bombers, 380 medium bombers, and 550 fighter bombers, one of the largest air attacks in World War II. Panzer Lehr was virtually wiped out, and the town was turned into rubble.'[22] The path was cleared for the allied advance, but many thousands of French citizens lost their lives in the process.

These actions were (and are still being) defended on the grounds that they were necessary to achieve military success. It is argued that the bombings were directly proportional to the military advantage gained. On the other hand, Walzer has robustly argued against the actions on moral grounds and in terms of just war.[23]

Another case of 'military necessity' was the practice of US troops in Korea and Vietnam of firing back at targets in areas laden with civilians when they came under enemy fire. When the troops were pinned down they automatically employed tanks to return fire into hillsides as well as call for air strikes and artillery support. These actions, whilst saving the lives of US troops, indiscriminately killed civilian men, women and children.[24]

Cohen points out that this was not an illegal act. The law of war is not necessarily moral; it 'allows troops under fire to fire back without ascertaining that there are no civilians mingled with the troops who are firing upon them. It allows troops under fire to fire back even if they know civilians are mingled with the enemy.'[25] Surely, this means that if soldiers are fired on, then lethally autonomous robots could be deployed in the same way as artillery or indiscriminate air strikes.

If countries at war or in conflict have armed autonomous robots that will save many of their soldiers' lives, will the deployment of those robots be deemed a military necessity? If it impacts both on the protection of soldiers' lives and on the ultimate success of the mission, then there will be a great temptation to use the technology. Imagine a situation where UAV deployment is what is giving state A the edge in an armed conflict with state B. Now imagine that state A has its communications disabled and its radio and GPS signals jammed. If state A can return to its advantageous position using its stock of (indiscriminate) autonomous armed UAVs to maintain its advantage, will it not do so?

Pushing the point home further, suppose that, having disrupted the remote control of state A's UAVs, state B now deploys autonomous attack craft; will state A not follow suit? Will concerns about keeping a person in the loop or unleashing possibly indiscriminate weapons prevent the use of lethal autonomous UAVs? It seems unlikely that a country will lose a war because it decides that moral superiority is more important than victory.

VI. CONCLUSION

After nearly a century of development, the Unmanned Aerial Vehicle has become perhaps the most desired asset amongst the modern militaries of the world. The military successes of UAVs in post-9/11 conflicts have created a rapid proliferation of the technology. Although there is currently a 'man-in-the-loop' for lethal targeting operations, that role will shrink incrementally until there is a capability for fully autonomous operation. The autonomous functions are likely to be ready long before robots will be able to distinguish between combatants and non-combatants in any way that requires battlefield awareness. They will not be able to reason appropriately or to make proportionality decisions barring some incredible and unpredictable technological breakthrough.

Concerns over the proliferation of the technology were expressed in this paper. The United States and Israel are currently well ahead of the field in terms of armed robot planes, but that may change in the near future. Russia has plans for armed unmanned combat aircraft, Iran claims to have them, and China is catching up quickly. More than 50 countries have been buying and developing the technology. It was pointed out that China will soon start selling and exporting its armed drones on the international market.

If they are not stopped, autonomous drones will likely be the tool of choice in future wars. While there is a lot of talk that humans will remain

in the loop to make lethal targeting decisions until the robots can be shown to be capable of obeying the principle of distinction, it is likely that military necessity will dictate that this constraint is dropped, ready or not. The mantra could then change to 'Don't worry, there will be technological fixes in the future'. This is similar to what the US has said about not signing the treaty banning cluster munitions.

The eventual proliferation of autonomous armed robots raises many questions of concern. What will happen when two complex algorithms meet in battle? Will the use of drones lead to lowering the bar to war because they appear to make it 'risk free'? Early indications from the current US administration are that drone use is already being considered an action short of warfare. What dangerous precedents are being set up by the current spate of CIA decapitations for when other countries have similar technology?

It is unclear what changes will need to be made to current international humanitarian law (IHL). IHL clearly covers the requirements for discrimination and proportionality. However, the mapping between the new technologies and IHL can be problematic when the operational detail of the new technology is not clear and keeps changing. Some clarifications and additions to IHL may be required, but it will be difficult to future-proof them. Armed autonomous robots are indiscriminate and disproportionate weapons and we must treat them as such now. As such the United Nations should consider placing them on the prohibited list of weapons under the Convention on Certain Conventional Weapons. We rely on possible future technological fixes at our peril.

NOTES

1. Teal group corporation website, http://bit.ly/psA7rB (accessed 1 September 2011).
2. I have personally read valid robotics reports for each of the following countries, and there may be many more: Australia, Austria, Brazil, Bulgaria, Canada, Chile, China, Columbia, Croatia, Czech Republic, Ecuador, Finland, France, Germany, Greece, Hungary, India, Indonesia, Iran, Israel, Italy, Japan, Jordan, Lebanon, Malaysia, Mexico, Netherlands, New Zealand, Norway, Pakistan, Peru, Philippines, Poland, Romania, Russia, Serbia, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, Tunisia, Turkey, United Arab Emirates, United Kingdom, USA, Vietnam.
3. 210 Report to Congress of the US-China Economic and Security Review Commission, November 2010, p 79, www.uscc.gov.
4. William Wan and Peter Finn, 'Global Race On to Match US Drone Capabilities' *Washington Post*, 4 July 2011, http://wapo.st/mfRa62 (accessed 10 August 2011).
5. War Powers Resolution of 1973, Public Law 93-148, 93rd Congress, HJ Res. 542, 7 November 1973, http://avalon.law.yale.edu/20th_century/warpower.asp (accessed 27 July 2011).
6. Letter from the President on the War Powers Resolution, 15 June 2011 www.whitehouse.gov/the-press-office/2011/06/15/letter-president-war-powers-resolution (accessed 17 August 2011).
7. Charlie Savage and Mark Landler (2011), 'White House Defends Continuing US Role in Libya Operation' *New York Times*, 15 June 2011, www.nytimes.com/2011/06/16/us/politics/16powers.html?pagewanted=all (accessed 3 August 2011).
8. Israel may have been using armed drones for longer but they denied this for several years despite eyewitness testimony. It cannot be verified here.
9. Nick Turse, 'Drone Surge: Today, Tomorrow and 2047' *Asia Times*, 26 January 2010, www.atimes.com/atimes/South_Asia/LA26Df01.html (accessed 17 July 2011).
10. Leon Panetta, *Director's Remarks at the Pacific Council on International Policy*, Central Intelligence Agency, 18 May 2009, www.cia.gov/news-information/speeches-testimony/directors-remarks-at-pacific-council.html (accessed 23 September 2010).
11. *Ibid.*
12. Harold Koh, 'The Obama Administration and International Law', speech delivered at the Annual Meeting of the American Society of International Law, Washington, DC, 25 March 2010.
13. Kenneth Anderson, Submission to US House of Representatives Committee on Oversight and Government Reform Subcommittee on National Security and Foreign Affairs, Subcommittee Hearing, *Drones II*, 28 April 2010, [20].
14. Philip Alston, Special Rapporteur on Extrajudicial, Summary or Arbitrary Executions, *Report of the Special Rapporteur on Extrajudicial, Summary or Arbitrary Executions, Addendum, Study on*

Targeted Killings (28 May 2010), UN Doc A/HRC/14/24/Add.6.

15. Noel Sharkey, 'Cassandra or the False Prophet of Doom: AI Robots and War' (2008) 23(4) *IEEE Intelligent Systems* 14.
16. 'USN Wants to Replace F-35s with UAVs' *Strategy Page* (online), 11 September 2011, www.strategypage.com/htmw/htnavai/articles/20110911.aspx (accessed 11 September 2011).
17. Mike Eisen, 'Amazon's $23,698,93 book about flies' (it is NOT junk, 22 April 2011), www.michaeleisen.org/blog/?p=358 (accessed 10 September 2011).
18. eg Sharkey (n 15); Noel Sharkey, 'Grounds for Discrimination: Autonomous Robot Weapons' (2008) 11(2) *RUSI Defence Systems* 86; Noel Sharkey, 'Saying No! to Lethal Autonomous Targeting' (2010) 9(4) *Journal of Military Ethics* 299.
19. But see also John S Ford, 'The Morality of Obliteration Bombing' (1944) 5 *Theological Studies* 261.
20. United Nations Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons Which May be Deemed to be Excessively Injurious or to Have Indiscriminate Effects, in force since 2 December 1983 and an annex to the Geneva Conventions of 12 August 1949.
21. ICJ, *Legality of the Threat or Use of Nuclear Weapons (General List No 95)* (8 July 1996).
22. SM Cohen, *Arms and Judgment: Law, Morality, and the Conduct of War in the Twentieth Century* (Westview, 1989) 34.
23. M Walzer, *Just and Unjust Wars* (Basic Books, 2006).
24. *Ibid.*
25. Cohen (n 22).

2.1.3 BIGGER MONSTER, WEAKER CHAINS

JAY STANLEY AND BARRY STEINHARDT

George Orwell's chilling novel *1984* depicts a society ruled by a totalitarian government bent on total manipulation and control of what people think and where every person is being constantly watched by Big Brother. When Orwell published his novel in 1949, the technologies that made such a dystopia possible were still in the realm of science fiction. But since that time, technological advances in digital cameras, computers, global positioning satellites, "data-mining," and other surveillance technologies have made "total information awareness" a real possibility. In addition, in response to the terrorist attacks of September 11, 2001, the U.S. government has increased its surveillance activities in the name of national security while keeping the details of what it has been doing secret from the general public. Civil libertarians, such as the authors of this report prepared for the American Civil Liberties Union in 2003, are alarmed at the erosion of personal privacy that these technological and political developments portend. Americans, they argue, are being increasingly watched, tracked, and recorded, both by private corporations and by a security-obsessed government. They fear that "privacy, while not dead, is on life support" and argue that the best way to prevent the emergence of a "surveillance society" is to enact strong new privacy laws based on the constitutional principle of the Fourth Amendment, laws that do not leave the right to privacy at the mercy of advancing technology.

Source: From *Bigger Monster, Weaker Chains: The Growth of an American Surveillance Society* by Jay Stanley and Barry Steinhardt, January 2005, American Civil Liberties Union Technology and Liberty Program. Reprinted by permission of the American Civil Liberties Union.

FOCUS QUESTIONS

1. What are some ways in which private corporations are increasingly using data surveillance to track and record our daily activities? Are you at all troubled by this trend toward the consolidation and commodification of personal information? Why or why not?
2. The authors claim that the post-9/11 legislation such as the USA PATRIOT Act expanded the government's authority to spy on its own citizens while weakening the system of judicial checks and balances on those powers. Explain why the authors think this is the case.
3. Why are the authors concerned about the possibility of a national ID card linked to a national database containing personal information about every citizen? Do you share their concerns, or do you believe that such a step would help the government protect U.S. citizens from criminals and terrorists? Where does the proper balance lie between personal privacy and national security?
4. Since this article was written, many new surveillance technologies have been developed and deployed. But there have also been some rulings, such as the recent decision by the U.S. Supreme Court (*United States v. Jones* 1/23/2012), which found that police placing a GPS tracking device on a suspect's vehicle constituted a search under the Fourth Amendment. Do you think that the laws protecting personal privacy are keeping up with new surveillance technology? Discuss.

KEYWORDS

biometrics, Carnivore, data profiling, data surveillance, digitization, Echelon, genetic information, GPS, PATRIOT Act, privacy, RFID chips, TIPS, video surveillance

THE GROWING SURVEILLANCE MONSTER

In the film *Minority Report*, which takes place in the United States in the year 2050, people called "Pre-cogs" can supposedly predict future crimes, and the nation has become a perfect surveillance society. The frightening thing is that except for the psychic Pre-cogs, the technologies of surveillance portrayed in the film already exist or are in the pipeline. Replace the Pre-cogs with "brain fingerprinting"—the supposed ability to ferret out dangerous tendencies by reading brain waves—and the film's entire vision no longer lies far in the future. Other new privacy invasions are coming at us from all directions, from video and data surveillance to DNA scanning to new data-gathering gadgets.

Video Surveillance

Surveillance video cameras are rapidly spreading throughout the public arena. A survey of surveillance cameras in Manhattan, for example, found that it is impossible to walk around the city without being recorded nearly every step of the way. And since September 11 the pace has quickened, with new cameras being placed not only in some of our most sacred public spaces, such as the National Mall in Washington and the Statue of Liberty in New York harbor, but on ordinary public streets all over America.

As common as video cameras have become, there are strong signs that, without public action, video surveillance may be on the verge of a revolutionary expansion in American life. There are three factors propelling this revolution:

1. *Improved technology.* Advances such as the digitization of video mean cheaper cameras, cheaper transmission of far-flung video feeds, and cheaper storage and retrieval of images.

2. *Centralized surveillance.* A new centralized surveillance center in Washington, DC, is an early indicator of what technology may bring. It allows officers to view images from video cameras across the city—public buildings and streets, neighborhoods, Metro stations, and even schools. With the flip of a switch, officers can zoom in on people from cameras a half-mile away.[1]

3. Unexamined assumptions that cameras provide security. In the wake of the September 11 attacks, many embraced surveillance as the way to prevent future attacks and prevent crime. But it is far from clear how cameras will increase security. U.S. government experts on security technology, noting that "monitoring video screens is both boring and mesmerizing," have found in experiments that after only 20 minutes of watching video monitors, "the attention of most individuals has degenerated to well below acceptable levels."[2] In addition, studies of cameras' effect on crime in Britain, where they have been extensively deployed, have found no conclusive evidence that they have reduced crime.[3]

These developments are creating powerful momentum toward pervasive video surveillance of our public spaces. If centralized video facilities are permitted in Washington and around the nation, it is inevitable that they will be expanded—not only in the number of cameras but also in their power and ability. It is easy to foresee inexpensive, one-dollar cameras being distributed throughout our cities and tied via wireless technology into a centralized police facility where the life of the city can be monitored. Those video signals could be stored indefinitely in digital form in giant but inexpensive databases, and called up with the click of a mouse at any time. With face recognition, the video records could even be indexed and searched based on who the systems identify—correctly, or, all too often, incorrectly.

Several airports around the nation, a handful of cities, and even the National Park Service at the Statue of Liberty have installed face recognition. While not nearly reliable enough to be effective as a security application,[4] such a system could still violate the privacy of a significant percentage of the citizens who appeared before it (as well as the privacy of those who do not appear before it but are falsely identified as having done so). Unlike, say, an iris scan, face recognition doesn't require the knowledge, consent, or participation of the subject; modern cameras can easily view faces from over 100 yards away.

Further possibilities for the expansion of video surveillance lie with unmanned aircraft, or drones, which have been used by the military and the CIA overseas for reconnaissance, surveillance, and targeting. Controlled from the ground, they can stay airborne for days at a time. Now there is talk of deploying them domestically. Senate Armed Services Committee Chairman John Warner (R, VA) said in December 2002 that he wants to explore their use in Homeland Security, and a number of domestic government agencies have expressed interest in deploying them. Drones are likely to be just one of many ways in which improving robotics technology will be applied to surveillance.[5]

The bottom line is that surveillance systems, once installed, rarely remain confined to their original purpose. Once the nation decides to go down the path of seeking security through video surveillance, the imperative to make it work will become overwhelming, and the monitoring of citizens in public places will quickly become pervasive.

Data Surveillance

An insidious new type of surveillance is becoming possible that is just as intrusive as video surveillance—what we might call "data surveillance." Data surveillance is *the collection of information about an identifiable individual, often from multiple sources, that can be assembled into a portrait of that person's activities.*[6] Most computers are programmed to automatically store and track usage data, and the spread of computer chips in our daily lives means that more and more of our activities leave behind "data trails." It will soon be possible to combine information from different sources to recreate an individual's activities with such detail that it becomes no different from being followed around all day by a detective with a video camera.

Some think comprehensive public tracking will make no difference, since life in public places is not "private" in the same way as life inside the home. This is wrong; such tracking would represent a radical change in American life. A woman who leaves her house, drives to a store, meets a friend for coffee, visits a museum, and then returns home may be in public all day, but her

life is still private in that she is the only one who has an overall view of how she spent her day. In America, she does not expect that her activities are being watched or tracked in any systematic way—she expects to be left alone. But if current trends continue, it will be impossible to have any contact with the outside world that is not watched and recorded.

The Commodification of Information

A major factor driving the trend toward data surveillance forward is the commodification of personal information by corporations. As computer technology exploded in recent decades, making it much easier to collect information about what Americans buy and do, companies came to realize that such data is often very valuable. The expense of marketing efforts gives businesses a strong incentive to know as much about consumers as possible so they can focus on the most likely new customers. Surveys, sweepstakes questionnaires, loyalty programs and detailed product registration forms have proliferated in American life—all aimed at gathering information about consumers. Today, any consumer activity that is *not* being tracked and recorded is increasingly being viewed by businesses as money left on the table.

On the Internet, where every mouse click can be recorded, the tracking and profiling of consumers is even more prevalent. Web sites can not only track what consumers buy, but what they *look at*—and for how long, and in what order. With the end of the Dot Com era, personal information has become an even more precious source of hard cash for those Internet ventures that survive. And of course Americans use the Internet not just as a shopping mall, but to research topics of interest, debate political issues, seek support for personal problems, and many other purposes that can generate deeply private information about their thoughts, interests, lifestyles, habits, and activities.

Genetic Privacy

The relentless commercialization of information has also led to the breakdown of some longstanding traditions, such as doctor-patient confidentiality. Citizens share some of their most intimate and embarrassing secrets with their doctors on the old-fashioned assumption that their conversations are confidential. Yet those details are routinely shared with insurance companies, researchers, marketers, and employers. An insurance trade organization called the Medical Information Bureau even keeps a centralized medical database with records on millions of patients. Weak new medical privacy rules will do little to stop this behavior.

An even greater threat to medical privacy is looming: genetic information. The increase in DNA analysis for medical testing, research, and other purposes will accelerate sharply in coming years and will increasingly be incorporated into routine health care.

Unlike other medical information, genetic data is a unique combination: both difficult to keep confidential and extremely revealing about us. DNA is very easy to acquire because we constantly slough off hair, saliva, skin cells and other samples of our DNA (household dust, for example, is made up primarily of dead human skin cells). That means that no matter how hard we strive to keep our genetic code private, we are always vulnerable to other parties' secretly testing samples of our DNA. The issue will be intensified by the development of cheap and efficient DNA chips capable of reading parts of our genetic sequences.

Already, it is possible to send away a DNA sample for analysis. A testing company called Genelex reports that it has amassed 50,000 DNA samples, many gathered surreptitiously for paternity testing. "You'd be amazed," the company's CEO told *U.S. News & World Report*. "Siblings have sent in mom's discarded Kleenex and wax from her hearing aid to resolve the family rumors."[7]

Not only is DNA easier to acquire than other medical information, revealing it can also have more profound consequences. Genetic markers are rapidly being identified for all sorts of genetic diseases, risk factors, and other characteristics. None of us knows what time bombs are lurking in our genomes.

The consequences of increased genetic transparency will likely include:

- *Discrimination by insurers.* Health and life insurance companies could collect DNA for use in deciding who to insure and what to charge them, with the result that a certain proportion of the population could become uninsurable. The insurance industry has already vigorously opposed efforts in Congress to pass meaningful genetic privacy and discrimination bills.
- *Employment discrimination.* Genetic workplace testing is already on the rise, and the courts have heard many cases. Employers desiring healthy, capable workers will always have an incentive to discriminate based on DNA—an incentive that will be even stronger as long as health insurance is provided through the workplace.
- *Genetic spying.* Cheap technology could allow everyone from schoolchildren to dating couples to nosy neighbors to routinely check out each other's genetic codes. A likely high-profile example: online posting of the genetic profiles of celebrities or politicians.

Financial Privacy

Like doctor-patient confidentiality, the tradition of privacy and discretion by financial institutions has also collapsed; financial companies today routinely put the details of their customers' financial lives up for sale.

A big part of the problem is the Gramm-Leach-Bliley Act passed by Congress in 1999. Although Gramm-Leach is sometimes described as a "financial privacy law," it created a very weak privacy standard—so weak, in fact, that far from protecting Americans' financial privacy, the law has had the effect of ratifying the increasing abandonment of customer privacy by financial companies.

Gramm-Leach effectively gives financial institutions permission to sell their customers' financial data to anyone they choose. That includes the date, amount, and recipient of credit card charges or checks a customer has written; account balances; and information about the flow of deposits and withdrawals through an account. Consumers provide a tremendous amount of information about themselves when they fill out applications to get a loan, buy insurance, or purchase securities, and companies can also share that information. In fact, the only information a financial company may NOT give out about you is your account number.

Under Gramm-Leach, you get no privacy unless you file complex paperwork, following a financial institution's precise instructions before a deadline they set, and repeating the process for each and every financial service provider who may have data about you. And it is a process that many companies intentionally make difficult and cumbersome; few let consumers "opt out" of data sharing through a Web site or phone number, or even provide a self-addressed envelope.

Gramm-Leach is an excellent example of the ways that privacy protections are being weakened even as the potential for privacy invasion grows.

New Data-Gathering Technologies

The discovery by businesses of the monetary value of personal information and the vast new project of tracking the habits of consumers has been made possible by advances in computers, databases and the Internet. In the near future, other new technologies will continue to fill out the mosaic of information it is possible to collect on every individual. Examples include the following:

- *Cell phone location data.* The government has mandated that manufacturers make cell phones capable of automatically reporting their location when an owner dials 911. Of course, those phones are capable of tracking their location at other times as well. And in applying the rules that protect the privacy of telephone records to this location data, the government is weakening those rules in a way that allows phone companies to collect and share data about the location and movements of their customers.
- *Biometrics.* Technologies that identify us by unique bodily attributes such as our fingerprints, faces, iris patterns, or DNA are already being proposed for inclusion on national ID cards and to identify airline passengers. Face recognition is spreading. Fingerprint scanners have

been introduced as security or payment mechanisms in office buildings, college campuses, grocery stores and even fast-food restaurants. And several companies are working on DNA chips that will be able to instantly identify individuals by the DNA we leave behind everywhere we go.

- *Black boxes.* All cars built today contain computers, and some of those computers are being programmed in ways that are not necessarily in the interest of owners. An increasing number of cars contain devices akin to the "black boxes" on aircraft that record details about a vehicle's operation and movement. Those devices can "tattle" on car owners to the police or insurance investigators. Already, one car rental agency tried to charge a customer for speeding after a GPS device in the car reported the transgression back to the company. And cars are just one example of how products and possessions can be programmed to spy and inform on their owners.
- *RFID chips.* RFID chips, which are already used in such applications as toll-booth speed passes, emit a short-range radio signal containing a unique code that identifies each chip. Once the cost of these chips falls to a few pennies each, plans are underway to affix them to products in stores, down to every can of soup and tube of toothpaste. They will allow everyday objects to "talk" to each other—or to anyone else who is listening. For example, they could let market researchers scan the contents of your purse or car from five feet away, or let police officers scan your identification when they pass you on the street.
- *Implantable GPS chips.* Computer chips that can record and broadcast their location have also been developed. In addition to practical uses such as building them into shipping containers, they can also serve as location "bugs" when, for example, hidden by a suspicious husband in a wife's purse. And they can be implanted under the skin (as can RFID chips).

If we do not act to reverse the current trend, data surveillance—like video surveillance—will allow corporations or the government to constantly monitor what individual Americans do every day. Data surveillance would cover *everyone*, with records of every transaction and activity squirreled away until they are sucked up by powerful search engines, whether as part of routine security checks, a general sweep for suspects in an unsolved crime, or a program of harassment against some future Martin Luther King.

Government Surveillance

Data surveillance is made possible by the growing ocean of privately collected personal data. But who would conduct that surveillance? There are certainly business incentives for doing so; companies called data aggregators (such as Acxiom and ChoicePoint) are in the business of compiling detailed databases on individuals and then selling that information to others. Although these companies are invisible to the average person, data aggregation is an enormous, multi-billion-dollar industry. Some databases are even "co-ops" where participants agree to contribute data about their customers in return for the ability to pull out cross-merchant profiles of customers' activities.

The biggest threat to privacy, however, comes from the government. Many Americans are naturally concerned about corporate surveillance, but only the government has the power to take away liberty—as has been demonstrated starkly by the post-September 11 detention of suspects without trial as "enemy combatants."

In addition, the government has unmatched power to centralize all the private sector data that is being generated. In fact, the distinction between government and private-sector privacy invasions is fading quickly. The Justice Department, for example, reportedly has an $8 million contract with data aggregator ChoicePoint that allows government agents to tap into the company's vast database of personal information on individuals.[8] Although the Privacy Act of 1974 banned the government from maintaining information on citizens who are not the targets of investigations, the FBI can now evade that requirement by simply purchasing information that has been collected by the private sector. Other proposals—such as the Pentagon's "Total Information Awareness" project and airline passenger profiling programs—would institutionalize

government access to consumer data in even more far-reaching ways.

Government Databases

The government's access to personal information begins with the thousands of databases it maintains on the lives of Americans and others. For instance:

- The FBI maintains a giant database that contains millions of records covering everything from criminal records to stolen boats and databases with millions of computerized fingerprints and DNA records.
- The Treasury Department runs a database that collects financial information reported to the government by thousands of banks and other financial institutions.
- A "new hires" database maintained by the Department of Health and Human Services, which contains the name, address, social security number, and quarterly wages of every working person in the U.S.
- The federal Department of Education maintains an enormous information bank holding years worth of educational records on individuals stretching from their primary school years through higher education. After September 11, Congress gave the FBI permission to access the database without probable cause.
- State departments of motor vehicles of course possess millions of up-to-date files containing a variety of personal data, including photographs of most adults living in the United States.

Communications Surveillance

The government also performs an increasing amount of eavesdropping on electronic communications. While technologies like telephone wiretapping have been around for decades, today's technologies cast a far broader net. The FBI's controversial "Carnivore" program, for example, is supposed to be used to tap into the e-mail traffic of a particular individual. Unlike a telephone wiretap, however, it doesn't cover just one device but (because of how the Internet is built) filters through *all* the traffic on the Internet Service Provider to which it has been attached. The only thing keeping the government from trolling through all this traffic is software instructions that are written by the government itself. (Despite that clear conflict of interest, the FBI has refused to allow independent inspection and oversight of the device's operation.)

Another example is the international eavesdropping program codenamed Echelon. Operated by a partnership consisting of the United States, Britain, Canada, Australia, and New Zealand, Echelon reportedly grabs e-mail, phone calls, and other electronic communications from its far-flung listening posts across most of the earth. (U.S. eavesdroppers are not supposed to listen in on the conversations of Americans, but the question about Echelon has always been whether the intelligence agencies of participating nations can set up reciprocal, back-scratching arrangements to spy on each others' citizens.) Like Carnivore, Echelon may be used against particular targets, but to do so its operators must sort through massive amounts of information about potentially millions of people. That is worlds away from the popular conception of the old wiretap where an FBI agent listens to one line. Not only the volume of intercepts but the potential for abuse is now exponentially higher.

The "Patriot" Act

The potential for the abuse of surveillance powers has also risen sharply due to a dramatic post-9/11 erosion of legal protections against government surveillance of citizens. Just six weeks after the September 11 attacks, a panicked Congress passed the "USA PATRIOT Act," an overnight revision of the nation's surveillance laws that vastly expanded the government's authority to spy on its own citizens and reduced checks and balances on those powers, such as judicial oversight. The government never demonstrated that restraints on surveillance had contributed to the attack, and indeed much of the new legislation had nothing to do with fighting terrorism. Rather, the bill represented a successful use of the terrorist attacks by the FBI to roll back unwanted checks on its

power. The most powerful provisions of the law allow for:

- *Easy access to records.* Under the PATRIOT Act, the FBI can force anyone to turn over records on their customers or clients, giving the government unchecked power to rifle through individuals' financial records, medical histories, Internet usage, travel patterns, or any other records. Some of the most invasive and disturbing uses permitted by the Act involve government access to citizens' reading habits from libraries and bookstores. The FBI does not have to show suspicion of a crime, can gag the recipient of a search order from disclosing the search to anyone, and is subject to no meaningful judicial oversight.
- *Expansion of the "pen register" exception in wiretap law.* The PATRIOT Act expands exceptions to the normal requirement for probable cause in wiretap law.[9] As with its new power to search records, the FBI need not show probable cause or even reasonable suspicion of criminal activity, and judicial oversight is essentially nil.
- *Expansion of the intelligence exception in wiretap law.* The PATRIOT Act also loosens the evidence needed by the government to justify an intelligence wiretap or physical search. Previously the law allowed exceptions to the Fourth Amendment for these kinds of searches only if "the purpose" of the search was to gather foreign intelligence. But the Act changes "the purpose" to "a significant purpose," which lets the government circumvent the Constitution's probable cause requirement even when its main goal is ordinary law enforcement.[10]
- *More secret searches.* Except in rare cases, the law has always required that the subject of a search be notified that a search is taking place. Such notice is a crucial check on the government's power because it forces the authorities to operate in the open and allows the subject of searches to challenge their validity in court. But the PATRIOT Act allows the government to conduct searches without notifying the subjects until long after the search has been executed.

Under these changes and other authorities asserted by the Bush Administration, U.S. intelligence agents could conduct a secret search of an American citizen's home, use evidence found there to declare him an "enemy combatant," and imprison him without trial. The courts would have no chance to review these decisions—indeed, they might never even find out about them.[11]

The "TIPS" Program

In the name of fighting terrorism, the Bush Administration has also proposed a program that would encourage citizens to spy on each other. The Administration initially planned to recruit people such as letter carriers and utility technicians, who, the White House said, are "well-positioned to recognize unusual events." In the face of fierce public criticism, the Administration scaled back the program, but continued to enlist workers involved in certain key industries. In November 2002 Congress included a provision in the Homeland Security Act prohibiting the Bush Administration from moving forward with TIPS.

Although Congress killed TIPS, the fact that the Administration would pursue such a program reveals a disturbing disconnect with American values and a disturbing lack of awareness of the history of governmental abuses of power. Dividing citizen from citizen by encouraging mutual suspicion and reporting to the government would dramatically increase the government's power by extending surveillance into every nook and cranny of American society. Such a strategy was central to the Soviet Union and other totalitarian regimes.

Loosened Domestic Spying Regulations

In May 2002, Attorney General John Ashcroft issued new guidelines on domestic spying that significantly increased the freedom of federal agents to conduct surveillance on American individuals and organizations. Under the new guidelines, FBI agents can infiltrate "any event that is open to the public," from public meetings and demonstrations to political conventions to church services to 12-step programs. This was the same basis upon which abuses were carried out by the

FBI in the 1950s and 1960s, including surveillance of political groups that disagreed with the government, anonymous letters sent to the spouses of targets to try to ruin their marriages, and the infamous campaign against Martin Luther King, who was investigated and harassed for decades. The new guidelines are purely for spying on Americans; there is a separate set of Foreign Guidelines that cover investigations inside the U.S. of foreign powers and terrorist organizations such as al Qaeda.

Like the TIPS program, Ashcroft's guidelines sow suspicion among citizens and extend the government's surveillance power into the capillaries of American life. It is not just the reality of government surveillance that chills free expression and the freedom that Americans enjoy. The same negative effects come when we are constantly forced to wonder whether we *might* be under observation—whether the person sitting next to us is secretly informing the government that we are "suspicious."

THE SYNERGIES OF SURVEILLANCE

Multiple surveillance techniques added together are greater than the sum of their parts. One example is face recognition, which combines the power of computerized software analysis, cameras, and databases to seek matches between facial images. But the real synergies of surveillance come into play with data collection.

The growing piles of data being collected on Americans represent an enormous invasion of privacy, but our privacy has actually been protected by the fact that all this information still remains scattered across many different databases. As a result, there exists a pent-up capacity for surveillance in American life today—a capacity that will be fully realized if the government, landlords, employers, or other powerful forces gain the ability to *draw together* all this information. A particular piece of data about you—such as the fact that you entered your office at 10:29 AM on July 5, 2001—is normally innocuous. But when enough pieces of that kind of data are assembled together, they add up to an extremely detailed and intrusive

Data Profiling and "Total Information Awareness"

Just how real this scenario is has been demonstrated by another ominous surveillance plan to emerge from the effort against terrorism: the Pentagon's "Total Information Awareness" program. The aim of this program is to give officials easy, unified access to every possible government and commercial database in the world.[12] According to program director John Poindexter, the program's goal is to develop "ultra-large-scale" database technologies with the goal of "treating the world-wide, distributed, legacy databases as if they were one centralized database." The program envisions a "full-coverage database containing all information relevant to identifying" potential terrorists and their supporters. As we have seen, the amount of available information is mushrooming by the day, and will soon be rich enough to reveal much of our lives.

The TIA program, which is run by the Defense Advanced Research Projects Agency (DARPA), not only seeks to bring together the oceans of data that are already being collected on people, but would be designed to afford what DARPA calls "easy future scaling" to embrace new sources of data as they become available. It would also incorporate other work being done by the military, such as their "Human Identification at a Distance" program, which seeks to allow identification and tracking of people from a distance, and therefore without their permission or knowledge.[13]

Although it has not received nearly as much media attention, a close cousin of TIA is also being created in the context of airline security. This plan involves the creation of a system for conducting background checks on individuals who wish to fly and then separating out either those who appear to be the most trustworthy passengers (proposals known as "trusted traveler") or flagging the least trustworthy (a proposal known as CAPS II, for Computer Assisted Passenger Screening) for special attention.

The *Washington Post* has reported that work is being done on CAPS II with the goal of creat-

to instantly pull together every passenger's travel history and living arrangements, plus a wealth of other personal and demographic information" in the hopes that the authorities will be able to "profile passenger activity and intuit obscure clues about potential threats." The government program would reportedly draw on enormous stores of personal information from data aggregators and other sources, including travel records, real estate histories, personal associations, credit card records, and telephone records. Plans call for using complex computer algorithms, including highly experimental technologies such as "neural networks," to sort through the reams of new personal information and identify "suspicious" people.[14]

The dubious premise of programs like TIA and CAPS II—that "terrorist patterns" can be ferreted out from the enormous mass of American lives, many of which will inevitably be quirky, eccentric, or riddled with suspicious coincidences—probably dooms them to failure. But failure is not likely to lead these programs to be shut down—instead, the government will begin feeding its computers more and more personal information in a vain effort to make the concept work. We will then have the worst of both worlds: poor security and a supercharged surveillance tool that would destroy Americans' privacy and threaten our freedom.

It is easy to imagine these systems being expanded in the future to share their risk assessments with other security systems. For example, CAPS could be linked to a photographic database and surveillance cameras equipped with face recognition software. Such a system might sound an alarm when a subject who has been designated as "suspicious" appears in public. The Suspicious Citizen could then be watched from a centralized video monitoring facility as he moves around the city.

In short, the government is working furiously to bring disparate sources of information about us together into one view, just as privacy advocates have been warning about for years. That would represent a radical branching off from the centuries-old Anglo-American tradition that the police conduct surveillance only where there is

would seek to protect us by monitoring *everyone* for signs of wrongdoing—in short, by instituting a giant dragnet capable of sifting through the personal lives of Americans in search of "suspicious" patterns. The potential for abuse of such a system is staggering.

The massive defense research capabilities of the United States have always involved the search for ways of outwardly defending our nation. Programs like TIA[15] involve turning those capabilities inward and applying them to the American people—something that should be done, if at all, only with extreme caution and plenty of public input, political debate, checks and balances, and Congressional oversight. So far, none of those things have been present with TIA or CAPS II.

National ID Cards

If Americans allow it, another convergence of surveillance technologies will probably center around a national ID card. A national ID would immediately combine new technologies such as biometrics and RFID chips along with an enormously powerful database (possibly distributed among the 50 states). Before long, it would become an overarching means of facilitating surveillance by allowing far-flung pools of information to be pulled together into a single, incredibly rich dossier or profile of our lives. Before long, office buildings, doctors' offices, gas stations, highway tolls, subways and buses would incorporate the ID card into their security or payment systems for greater efficiency, and data that is currently scattered and disconnected will get organized around the ID and lead to the creation of what amounts to a national database of sensitive information about American citizens.

History has shown that databases created for one purpose are almost inevitably expanded to other uses; Social Security, which was prohibited by federal law from being used as an identifier when it was first created, is a prime example. Over time, a national ID database would inevitably contain a wider and wider range of information and become accessible to more and more

further and further removed from its original justification.

The most likely route to a national ID is through our driver's licenses. Since September 11, the American Association of Motor Vehicle Administrators has been forcefully lobbying Congress for funds to establish nationwide uniformity in the design and content of driver's licenses—and more importantly, for tightly interconnecting the databases that lie behind the physical licenses themselves.

An attempt to retrofit driver's licenses into national ID cards will launch a predictable series of events bringing us toward a surveillance society:

- Proponents will promise that the IDs will be implemented in limited ways that won't devastate privacy and other liberties.
- Once a limited version of the proposals is put in place, its limits as an anti-terrorism measure will quickly become apparent. Like a dam built halfway across a river, the IDs cannot possibly be effective unless their coverage is total.
- The scheme's ineffectiveness—starkly demonstrated, perhaps, by a new terrorist attack—will create an overwhelming imperative to "fix" and "complete" it, which will turn it into the totalitarian tool that proponents promised it would never become.

A perfect example of that dynamic is the requirement that travelers present driver's licenses when boarding airplanes, instituted after the explosion (now believed to have been mechanical in cause) that brought down TWA Flight 800 in 1996. On its own, the requirement was meaningless as a security measure, but after September 11 its existence quickly led to calls to begin tracking and identifying citizens on the theory that "we already have to show ID, we might as well make it mean something."

Once in place, it is easy to imagine how national IDs could be combined with an RFID chip to allow for convenient, at-a-distance verification of ID. The IDs could then be tied to access control points around our public places, so that the unauthorized could be kept out of office buildings, apartments, public transit, and secure public buildings. Citizens with criminal records, poor CAPS ratings or low incomes could be barred from accessing airports, sports arenas, stores, or other facilities. Retailers might add RFID readers to find out exactly who is browsing their aisles, gawking at their window displays from the sidewalk or passing by without looking. A network of automated RFID listening posts on the sidewalks and roads could even reveal the location of all citizens at all times. Pocket ID readers could be used by FBI agents to sweep up the identities of everyone at a political meeting, protest march, or Islamic prayer service.

CONCLUSION

If we do not take steps to control and regulate surveillance to bring it into conformity with our values, we will find ourselves being tracked, analyzed, profiled, and flagged in our daily lives to a degree we can scarcely imagine today. We will be forced into an impossible struggle to conform to the letter of every rule, law, and guideline, lest we create ammunition for enemies in the government or elsewhere. Our transgressions will become permanent Scarlet Letters that follow us throughout our lives, visible to all and used by the government, landlords, employers, insurance companies and other powerful parties to increase their leverage over average people. Americans will not be able to engage in political protest or go about their daily lives without the constant awareness that we are—or could be—under surveillance. We will be forced to constantly ask of even the smallest action taken in public, "Will this make me look suspicious? Will this hurt my chances for future employment? Will this reduce my ability to get insurance?" The exercise of free speech will be chilled as Americans become conscious that their every word may be reported to the government by FBI infiltrators, suspicious fellow citizens or an Internet Service Provider.

Many well-known commentators like Sun Microsystems CEO Scott McNealy have already pronounced privacy dead. The truth is that a surveillance society does loom over us, and privacy, while not yet dead, is on life support.

Heroic measures are required to save it.

Four main goals need to be attained to prevent this dark potential from being realized: a change in the terms of the debate, passage of comprehensive privacy laws, passage of new laws to regulate the powerful and invasive new technologies that have and will continue to appear, and a revival of the Fourth Amendment to the U.S. Constitution.

1. Changing the Terms of the Debate

In the public debates over every new surveillance technology, the forest too often gets lost for the trees, and we lose sight of the larger trend: the seemingly inexorable movement toward a surveillance society. It will always be important to understand and publicly debate every new technology and every new technique for spying on people. But unless each new development is also understood as just one piece of the larger surveillance mosaic that is rapidly being constructed around us, Americans are not likely to get excited about a given incremental loss of privacy like the tracking of cars through toll booths or the growing practice of tracking consumers' supermarket purchases.

We are being confronted with fundamental choices about what sort of society we want to live in. But unless the terms of the debate are changed to focus on the forest instead of individual trees, too many Americans will never even recognize the choice we face, and a decision against preserving privacy will be made by default.

2. Comprehensive Privacy Laws

Although broad-based protections against government surveillance, such as the wiretap laws, are being weakened, at least they exist. But surveillance is increasingly being carried out by the private sector—frequently at the behest of government—and the laws protecting Americans against non-governmental privacy invasions are pitifully weak.

In contrast to the rest of the developed world, the U.S. has no strong, comprehensive law protecting privacy—only a patchwork of largely inadequate protections. For example, as a result of many legislators' discomfort over the disclosure of Judge Robert Bork's video rental choices during his Supreme Court confirmation battle, video records are now protected by a strong privacy law. Medical records are governed by a separate, far weaker law that allows for widespread access to extremely personal information. Financial data is governed by yet another "privacy" law— Gramm-Leach—which as we have seen really amounts to a license to share financial information. Another law protects only the privacy of children under age 13 on the Internet. And layered on top of this sectoral approach to privacy by the federal government is a geographical patchwork of constitutional and statutory privacy protections in the states.

The patchwork approach to privacy is grossly inadequate. As invasive practices grow, Americans will face constant uncertainty about when and how these complex laws protect them, contributing to a pervasive sense of insecurity. With the glaring exception of the United States, every advanced industrialized nation in the world has enacted overarching privacy laws that protect citizens against private-sector abuses. When it comes to this fundamental human value, the U.S. is an outlaw nation. For example, the European Union bars companies from evading privacy rules by transferring personal information to other nations whose data-protection policies are "inadequate." That is the kind of law that is usually applied to Third World countries, but the EU counts the United States in this category.

We need to develop a baseline of simple and clear privacy protections that crosses all sectors of our lives and give it the force of law. Only then can Americans act with a confident knowledge of when they can and cannot be monitored.

3. New Technologies and New Laws

The technologies of surveillance are developing at the speed of light, but the body of law that protects us is stuck back in the Stone Age. In the past, new technologies that threatened our privacy, such as telephone wiretapping, were assimilated over time into our society. The legal system had time to adapt and reinterpret existing laws, the political system had time to consider and enact

new laws or regulations, and the culture had time to absorb the implications of the new technology for daily life. Today, however, change is happening so fast that none of this adaptation has time to take place—a problem that is being intensified by the scramble to enact unexamined anti-terrorism measures. The result is a significant danger that surveillance practices will become entrenched in American life that would never be accepted if we had more time to digest them.

Since a comprehensive privacy law may never be passed in the U.S.—and certainly not in the near future—law and legal principles must be developed or adapted to rein in particular new technologies such as surveillance cameras, location-tracking devices, and biometrics. Surveillance cameras, for example, must be subject to force-of-law rules covering important details like when they will be used, how long images will be stored, and when and with whom they will be shared.

4. Reviving the Fourth Amendment

The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no warrants shall issue, but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

FOURTH AMENDMENT
TO THE U.S. CONSTITUTION

The Fourth Amendment, the primary Constitutional bulwark against Government invasion of our privacy, was a direct response to the British authorities' use of "general warrants" to conduct broad searches of the rebellious colonists.

Historically, the courts have been slow to adapt the Fourth Amendment to the realities of developing technologies. It took almost 40 years for the U.S. Supreme Court to recognize that the Constitution applies to the wiretapping of telephone conversations.[16]

In recent years—in no small part as the result of the failed "war on drugs"—Fourth Amendment principles have been steadily eroding. The circumstances under which police and other government officials may conduct warrantless searches have been rapidly expanding. The courts have allowed for increased surveillance and searches on the nation's highways and at our "borders" (the legal definition of which actually extends hundreds of miles inland from the actual border). And despite the Constitution's plain language covering "persons" and "effects," the courts have increasingly allowed for warrantless searches when we are outside of our homes and "in public." Here the courts have increasingly found we have no "reasonable expectation" of privacy and that therefore the Fourth Amendment does not apply.

But like other Constitutional provisions, the Fourth Amendment needs to be understood in contemporary terms. New technologies are endowing the government with the 21st century equivalent of Superman's X-ray vision. Using everything from powerful video technologies that can literally see in the dark, to biometric identification techniques like face recognition, to "brain fingerprinting" that can purportedly read our thoughts, the government is now capable of conducting broad searches of our "persons and effects" while we are going about our daily lives—even while we are in "public."

The Fourth Amendment is in desperate need of a revival. The reasonable expectation of privacy cannot be defined by the power that technology affords the government to spy on us. Since that power is increasingly limitless, the "reasonable expectation" standard will leave our privacy dead indeed.

But all is not yet lost. There is some reason for hope. In an important pre-9/11 case, *Kyllo vs. U.S.*,[17] the Supreme Court held that the reasonable expectation of privacy could not be determined by the power of new technologies. In a remarkable opinion written by conservative Justice Antonin Scalia, the Court held that without a warrant the police could not use a new thermal imaging device that searches for heat sources to conduct what was the functional equivalent of a warrantless search for marijuana cultivation in Danny Kyllo's home.

The Court specifically declined to leave Kyllo "at the mercy of advancing technology." While *Kyllo* involved a search of a home, it enunciates an important principle: the Fourth Amendment must adapt to new technologies. That principle can and should be expanded to general use. The Framers never expected the Constitution to be read exclusively in terms of the circumstances of 1791.

NOTES

1. Jess Bravin, "Washington Police to Play 'I Spy' with Cameras, Raising Concerns," *Wall Street Journal*, Feb. 13, 2002.
2. See http://www.ncjrs.org/school/ch2a5.html.
3. See http://www.scotcrim.u-net.com/researchc2.htm.
4. The success rate of face recognition technology has been dismal. The many independent findings to that effect include a trial conducted by the U.S. military in 2002, which found that with a reasonably low false-positive rate, the technology had less than a 20% chance of successfully identifying a person in its database who appeared before the camera. See http://www.aclu.org/issues/privacy/FINAL1 Final Steve King.pdf, 17th slide.
5. Richard H. P. Sia, "Pilotless Aircraft Makers Seek Role for Domestic Uses," *CongressDaily*, Dec. 17, 2002.
6. Data surveillance is often loosely referred to as "data mining." Strictly speaking, however, data mining refers to the search for hidden patterns in large, pre-existing collections of data (such as the finding that sales of both beer and diapers rise on Friday nights). Data mining need not involve personally identifiable information. Data surveillance, on the other hand, involves the collection of information about an identifiable individual. Note, however, that when data surveillance is carried out on a mass scale, a search for patterns in people's activities—data mining—can then be conducted as well. This is what appears to be contemplated in the Total Information Awareness and CAPS II programs (see below).
7. Dana Hawkins, "As DNA Banks Quietly Multiply, Who Is Guarding the Safe?" *U.S. News & World Report*, Dec. 2, 2002.
8. Glenn R. Simpson, "Big Brother-in-Law: If the FBI Hopes to Get the Goods on You, It May Ask ChoicePoint," *Wall St. Journal*, Apr. 13, 2001.
9. The expanded exception involves what are called "pen register/trap & trace" warrants that collect "addressing information" but not the content of a communication. Those searches are named after devices that were used on telephones to show a list of telephone numbers dialed and received (as opposed to tapping into actual conversations). The PATRIOT Act expands the pen register exception onto the Internet in ways that will probably be used by the government to collect the actual content of communications and that allow nonspecific "nationwide" warrants in violation of the Fourth Amendment's explicit requirement that warrants "must specify the place to be searched."
10. In August, the secret "FISA" court that oversees domestic intelligence spying released an opinion rejecting a Bush Administration attempt to allow criminal prosecutors to use intelligence warrants to evade the Fourth Amendment entirely. The court noted that agents applying for warrants had regularly filed false and misleading information. In November 2002, however, the FISA appeals court (three judges chosen by Supreme Court Chief Justice William Rehnquist), meeting for the first time ever, ruled in favor of the government.
11. See Charles Lane, "In Terror War, 2nd Track for Suspects," *Washington Post*, Dec. 1, 2002. Online at http://www.washingtonpost.com/wp-dyn/articles/A58308-2002Nov30.html.
12. See "Pentagon Plans a Computer System That Would Peek at Personal Data of Americans," *New York Times*, Nov. 9, 2002, "US Hopes to Check Computers Globally," *Washington Post*, Nov. 12, 2002, "The Poindexter Plan," *National Journal*, Sept. 7, 2002.
13. Quotes are from the TIA homepage at http://www.darpa.mil/iao/index.htm and from public 8/2/02 remarks by Poindexter, online at http://www.fas.org/irp/agency/dod/poindexter.html.
14. Robert O'Harrow Jr., "Intricate Screening of Fliers in Works," *Washington Post*, Feb. 1, 2002, p. A1.
15. The TIA is just one part of a larger post-9/11 expansion of federal research and development efforts. The budget for military R&D spending alone has been increased by 18% in the current fiscal year to a record $58.8 billion. Bob Davis, "Massive Federal R&D Initiative to Fight Terror Is Under Way," *Wall Street Journal*, Nov. 25, 2002.
16. In 1967 the Supreme Court finally recognized the right to privacy in telephone conversations in the case *Katz v. U.S.* (389 US 347), reversing the 1928 opinion *Olmstead v. U.S.* (277 US 438).
17. 190 F.3d 1041, 2001.

2.1.4 IN PRAISE OF BIG BROTHER

JAMES STACEY TAYLOR

In Selection 2.1.3, "Bigger Monster, Weaker Chains," Jay Stanley and Barry Steinhardt decry the growing use of surveillance technologies that they see as a threat to privacy. Not everyone agrees, however, that increased surveillance by governments and corporations is such a bad thing. James Stacey Taylor presents a contrarian view in which he argues that what matters is *not* whether someone is watching over you but *what* he or she does with the information obtained by contemporary surveillance technologies. Rather than opposing the use of these technologies for gathering personal information, we should encourage their use as long as we observe the distinction between the surveillance itself and the question of who is permitted to access the gathered information and the ways in which that information is used. This argument is based on the principle of polypotency of technologies under which the same technology can be used for good or bad purposes. Stacey's argument turns on our accepting the idea that, with proper legal safeguards against governmental abuse, omnipresent surveillance can actually promote important social benefits such as deterring crime and lessening the need to subpoena reluctant witnesses to testify in trials. He argues that, rather than violating rights, the morally legitimate use of surveillance technologies can actually make us safer and enhance our autonomy and privacy.

FOCUS QUESTIONS

1. What is Taylor's argument for thinking that use of surveillance technologies by the state should be morally acceptable to both rights theorists and consequentialists?
2. Under what circumstances should the state be permitted to access the personal information that it gathers through the use of constant surveillance? What limitations and safeguards does Taylor propose to minimize abuse of these powers?
3. Does it matter what kind of government employs surveillance technology to monitor personal behavior? Should U.S. citizens be wary of increased reliance on these technologies by their own government?
4. Can a person be harmed by being spied upon if he or she does not know that this is happening? Do you agree with Taylor's view? Explain.
5. Comparing the argument in this selection with Selection 2.1.3, which do you think has the better argument? Discuss.

KEYWORDS

autonomy, Big Brother, consequentialism, power, privacy, Spyware, subpoena, surveillance,

Source: From "In Praise of Big Brother: Why We Should Learn to Stop Worrying and Love (Some) Government Surveillance," by James Stacey Taylor. Revised version. Original version published in *Public Affairs Quarterly* 19, no. 3 (2005): 227–246. Reprinted by permission.

In recent years surveillance technology has undergone a revolution. Spy satellites are now so accurate that they can be used to track the movements of individual people, and even read license plates on cars. It is now simple to intercept faxes, pager messages, and telephone messages, for rooms to be bugged, and for tracking devices to be installed on vehicles, goods, and even individual persons. Spyware software can now be covertly and remotely implanted onto Internet-linked computers to monitor the keystrokes that their users make and the websites that they visit. And in many areas of the developed world persons now live much of their public lives under the panoptic gaze of closed-circuit television monitors that record their every move.

This revolution in surveillance technology is often regarded with horror. Discussions of its ethical implications frequently draw analogies between it and Big Brother's sinister surveillance of the citizens of Oceania in George Orwell's *1984*. Worse yet, as many of the participants in such discussions note, the surveillance capabilities that are now available to governments and corporations dwarf that which Big Brother had access to. In *1984*, people could be sure that they could not be watched by Big Brother's telescreens if they were in a crowd.[1] Face-recognition software now renders such a hope futile. In *1984* people could escape the omnipresence of Big Brother's telescreens by going out of the cities and into the country, where they only had to take care that their conversations were not monitored by hidden microphones.[2] Such escape is now impossible, for spy satellites can be used to monitor people wherever they go.

With the dystopic vision of *1984* lurking in the background, it is clear that one need not possess any pronounced Luddite tendencies to oppose the expanding use of surveillance technology. It will be argued in this paper, however, that rather than opposing such an expansion of surveillance technology, its use should be *encouraged*—and not only in the public realm. Indeed, the State should place all of its citizens under surveillance at all times and in all places, including their offices,

AN OVERVIEW OF THE ARGUMENT

At first sight, the conclusion that the State should place all of its citizens under constant surveillance is alarming. Yet this alarm can be dispelled once it is realized that this conclusion flows naturally from the plausible and widely held view that, in certain circumstances, it is morally permissible for the State to secure information about past events. It is, for example, widely held that in certain circumstances it is morally permissible for the State to compel witnesses to testify about past events in criminal trials. The State, however, can only use hindsight to determine what information it is morally permitted to have access to, for it will only become clear in retrospect what information is relevant to (for example) solving a crime or judging mitigating circumstances. To ensure that it gleans all of the information that it is morally permitted to access, then, the State can gather information about all events that occur, provided it only accesses that which it is morally permitted. Given this, then, the State is in principle morally permitted to place its citizens under constant surveillance. The purpose of this paper is not, however, only to show that the State is morally permitted to place its citizens under constant surveillance. It is also to show that a situation in which the State used such surveillance would be morally *preferable* to one where it did not. To establish this it will be shown that under such a system of surveillance it is likely that crime will decrease, that justice will be better served, and that fewer costs will be imposed on witnesses. Finally, four objections to the use of such a system of State surveillance will be examined and rejected: that it would be open to abuse, that it is not morally permissible for the State to secure information about past events, and that such surveillance would violate citizens' privacy, or autonomy, or both.

Before moving to develop the above argument in favor of constant State surveillance an important initial clarification is in order. This argument is based on two claims: (i) that if it is ever morally permissible for the State to secure information about past events, then it is morally

surveillance devices, and (ii) that in some cases it is permissible for the State to secure information about past events. As it stands, however, it is not clear whether (ii) is the claim that the State's acquisition of information respects a person's moral rights (e.g., to privacy or autonomy), or the claim that in some cases the consequences of securing such information justify the State's acquisition of it. That is, it is not clear whether this argument is one that would be acceptable to a rights theorist, or only to a consequentialist.[3] This unclarity can be dispelled once it is recognized that the argument presented here is compatible with *both* a rights-based and a consequentialist approach to ethics.[4] This argument is acceptable to a rights theorist because, as will be argued below, when the State is morally permitted to access information it will not violate anyone's rights in so doing. If the arguments below are sound, then claim (ii) will be true if one accepts a rights-based approach to ethics. This argument is also acceptable to consequentialists, who could also accept claim (ii).[5] For a consequentialist to accept (ii) it must be the case that there is at least *one* time at which it is clear that the consequences justify the State's acquisition of information about a certain past event. And, as will be argued below, such a situation is simple to envisage. As such, the argument of this paper is amenable to both rights theorists and consequentialists alike.[6]

SURVEILLANCE AND TESTIMONY

Despite appearances, then, the conclusion that there is no principled reason why a State should not place all of its citizens under constant surveillance is by no means a radical one. It is merely an extrapolation of the widely accepted view that, under certain circumstances, it is morally permissible for agents of the State to secure information about past events. It is, for example, widely accepted that it is morally permissible for judges to subpoena witnesses to require them to disclose information. It is also widely accepted that it is morally permissible for judges to permit law enforcement agencies to install surveillance devices to monitor the activities of persons they suspect of criminal activity, provided that such agencies demonstrate that they have probable cause for their suspicions.[7] The claim that it is morally permissible for agents of the State to secure information about past events is thus innocuous.[8] Moreover, if it is morally permissible for agents of the State to secure information about past events through the subpoenaing of witnesses or the use of surveillance devices, then there should be no moral bar to their gaining such information through securing access to records of past events that might have been generated by preexisting surveillance devices. (Provided, of course, that in securing such information the agents of the State are subject to the same restrictions on their powers that are in place with respect to their subpoenaing of witnesses or the placing of suspects under surveillance.)

However, that it is morally permissible for agents of the State to secure access to information about past events that has been gathered through the use of surveillance devices does not show that it is morally permissible for the State to place its citizens under constant surveillance. To reach this further conclusion it must first be noted that it cannot (typically) be known in advance what information it would be morally permissible for the agents of the State to have access to. Instead, one can only use hindsight to determine what events it would be morally permissible for them to access information about. For example, it could not have been predicted that knowledge of the actions performed by Billy Nolan Lovelady, a worker at the Texas Square Book Depository in Dallas on 22nd November, 1963, would be relevant to assessing who assassinated the president of the United States. However, the defenders of Lee Harvey Oswald (who was accused of this assassination) claimed that at the time of the shooting Oswald was standing unarmed in the doorway of the Texas Square Book Depository and produced a photograph purporting to prove this. Once this had occurred, knowledge of Lovelady's actions became relevant in showing that the man in the photograph who was claimed to be Oswald was actually Lovelady. In this situation, then, it was morally permissible for agents of the State to take steps to secure information about Lovelady's

whereabouts at the time in question to disprove Oswald's alibi—and this could not have been predicted beforehand.

As well as it being clear that one usually cannot know in advance what information the agents of the State are morally permitted to secure, it is also clear that once it becomes evident what information they are permitted to secure, gathering this is frequently difficult—as the continuing controversy over "who shot JFK?" indicates. Putting these two facts together, then, it is evident that there should be no moral bar against the agents of the State installing surveillance devices to secure all the information that they might potentially be permitted to secure access to, provided that *they only access the information thus gathered in circumstances where this access is morally permissible.* Thus, since it is *potentially* morally permissible for agents of the State to secure information about events no matter when or where they take place, it is morally permissible for the State to subject its citizens to surveillance at all times and in all places, provided that the information thus gathered is only accessed in the morally appropriate circumstances.

FURTHER CLARIFYING THE ARGUMENT

Yet despite the above disclaimer that this argument is by no means as radical as it appears the conclusion that it is morally permissible for the State to subject its citizens to constant surveillance might still seem chilling. To alleviate this concern the above pro-surveillance argument should be clarified in two ways. First, it must be emphasized that although this argument leads to the conclusion that it is morally permissible for the State to place its citizens under constant surveillance, it does *not* lead to the conclusion that it is morally permissible for the State to have access to *all* of the information that its surveillance devices secure. Instead, this argument leads only to the conclusion that the State should have access to the information recorded by these devices that it is *morally permissible* for it to have access to. Of course, the question of the extent to which the State is morally permitted to secure information about the actions of its citizens, or the events that they are involved in, is a vexed one. It is, for example, debatable as to how far the privilege that is accorded to information divulged within certain professional relationships (such as those between clergy and penitents,[9] doctors and patients,[10] lawyers and clients,[11] counselors and clients,[12] or journalists and sources[13]) should extend, and for what reasons (if any) such privilege should be trumped by other considerations.[14] It is also debatable as to whether the increased powers accorded to law enforcement agencies by the Patriot Act to secure information about suspects or "persons of interest" are morally legitimate.[15] However, it must be stressed that the question of *what* information the State that subjects its citizens to constant and universal surveillance should have access to is separate from the question of *whether it is morally permissible* to *place them under surveillance at all*—and it is the *latter* question that is at issue here.

The second way in which this pro-surveillance argument must be clarified is related to the first. In cases where agents of the State *are* morally permitted to access information about the actions of others the information that they are morally permitted to secure is limited to the *minimum* that is needed for them to achieve the legitimate purposes for which they need it. Accordingly they are not morally permitted to have access to, for example, videotaped records of the actions of others if the information that they are morally permitted to secure could be gleaned from an automatic transcription of such videotape. Thus, although it might be morally permissible for an agent of the State to secure detailed information about a criminal defendant's past adultery if this is relevant to the State's case against him in a criminal trial, it does not follow that it is also morally permissible for this agent to watch a videotape of the defendant in the act. Similarly, the agents of the State would only be morally permitted to place citizens under surveillance in private areas such as their homes or offices if they could achieve this without accessing more information than they were morally permitted to access. It would thus not be morally permissible, for example, for the

install surveillance devices in it. This is because in doing so they would access more information than they were morally permitted to access (e.g., they would find out information about his domestic habits and the interior of his home during the installation process) for the installation of such devices would have to occur *before* there were reasons to believe that the State was morally permitted to secure such information to detect and prosecute criminal activity. Indeed, it is likely that, with respect to the interiors of the homes of most of the State's citizens, such reasons will never be forthcoming.

This second clarification of the above pro-surveillance argument is important in three respects. First, it underscores the fact that this argument does not cede to the State any information-gathering powers that it does not already enjoy. Second, it emphasizes that the conclusion of this argument is a parsimonious one, insofar as it only supports the State's securing the *minimum* amount of information that it needs for its morally legitimate purposes. This argument will thus *not* support the use of any surveillance devices that enable the State to secure more information than is necessary for its morally legitimate purposes. As such, although this argument does show that the State is morally permitted to use surveillance devices to gather information about events that occur in private areas (such as private offices and houses) it restricts its use of such devices to those whose use does not necessitate the State's securing any more information than it is morally permitted to secure. Thus, as noted above, this argument will not support the installation of surveillance devices in, for example, persons' homes without their consent. Although it does support the use of any technology, such as, for example, infrared thermal imaging systems that record the movements of persons in their homes, which could secure information about the events that occurred to them, that would not entail the State's securing of information that it was not entitled to.[16] Similarly, this argument will also support the use of future surveillance technology that can similarly be used to record events that occur in persons' homes without having to enter them. Finally, and in a related vein, once it becomes possible for surveillance technology automatically to provide a written narrative of events, the use of any type of surveillance technology that provides *more* than this information, such as, for example, manned closed-circuit television monitors and videotapes, will *cease* to be morally justified by the above pro-surveillance argument. Rather than justifying the *expansion* of the State's powers of surveillance, then, the above pro-surveillance argument will instead eventually justify their *curtailment*.

THE ADVANTAGES OF UNIVERSAL STATE SURVEILLANCE

If the above pro-surveillance argument is sound, then it is morally permissible for a State to subject its citizens to constant surveillance. Given the horror with which the current proliferation of surveillance devices is usually greeted, this conclusion is not likely to be a popular one. However, before turning to rebut the objections that it will be faced with, some of its advantages should be outlined.

The most obvious advantages to the State's installing such a surveillance system would result from the fact that witnesses would no longer be needed in either criminal or civil cases, for their testimony would be supplanted by information supplied by surveillance devices. Unlike witness testimony, this information would be accurate. It would, for example, be unaffected by any biases (whether conscious or unconscious) that human witnesses might be subject to and which could taint their testimony. It would also be free from distortions, whether deliberate (e.g., the witness is lying, or omitting parts of the truth) or accidental (e.g., the witness has a faulty memory). Moreover, the juries and judges to whom the information taken from surveillance devices would be presented in a court could take it at its face value, rather than having subjectively to assess its accuracy in the light of the perceived reliability of the witness from whom it was taken. Furthermore, the fact that witnesses would no longer be needed under a system of constant and universal State surveillance would also benefit those who would otherwise have served in this capacity. Most obviously, they would no longer be burdened with the

task of testifying, which might have required them to travel, or to take time off work. They would also be relieved from any threats that they might have faced from persons who would be adversely affected either by their testimony (e.g., the defendants, or their associates) or by its lack (e.g., the prosecutors).

A system of constant State surveillance would have other advantages, too. Under the current criminal justice system, a wealthy defendant who is innocent of the charges that she is faced with can use her wealth to hire private investigators to demonstrate her innocence, either by finding persons who witnessed the crime of which she is accused or by finding persons who can provide her with a legitimate alibi. This option is not open to poorer defendants who are similarly innocent, but who cannot afford to hire private investigators. Since this is so, innocent, poor defendants are more likely than innocent, wealthy defendants to accept plea-bargains, or to be convicted of crimes that they did not commit. If, however, a poor person were to be accused of a crime in a State that subjected its citizens to constant surveillance, the judge in her case would be morally justified (indeed, would be morally *required*) in enabling the defense to secure information that would prove her innocence, and that would have been gathered by the State's surveillance devices. A State's use of constant surveillance could thus reduce the number of persons who are wrongfully convicted. This would not only be good in itself, but it would also lead to a more equitable justice system, for the disparity in wrongful conviction rates between the wealthy, who could use their wealth to prove their innocence, and the poor could be eliminated.

In addition to these advantages constant State surveillance would also benefit the State's citizenry at large, by serving as a deterrent to crime. This would not be because the citizens who lived under State surveillance would never know when they were being watched, and so would refrain from committing crimes for fear of being caught in the act, as would persons in both *1984* and Jeremy Bentham's Panopticon.[17] As is clear from the second clarification of this pro-surveillance argument, above, this argument does not justify the State watching its citizens, for in doing so it would acquire more information about them than it is morally permitted to acquire. However, were the State to subject its citizens to the degree of surveillance that would be justified on the above argument, its citizens would know that whenever a crime was committed its performance would be recorded, and so its perpetrator would be likely to be apprehended. This knowledge would deter many potential criminals from committing crimes.[18] It would also deter law enforcement officers from extending the limits of their authority.[19] To be sure, the knowledge that one's criminal act would be recorded would not deter all potential criminals from committing crimes. Some persons would still commit crimes of passion, and of anger; others would succumb to the temptation to commit an opportunistic theft, act under the influence of drugs or alcohol, or act on the belief that they would not be caught. But most persons would simply judge the commission of crime not worthwhile once they realized that they were subject to constant State surveillance, and so such surveillance would serve as an effective deterrent to most criminal activity.

TWO INITIAL OBJECTIONS

Despite the advantages that would accrue to a system of constant State surveillance of the type that is defended above, there are two immediate objections that its proponents must face. The first is that such a surveillance system would be open to abuse. The second denies the claim that it is morally permissible for agents of the State (e.g., judges) to compel witness testimony or authorize information-gathering surveillance.

The Objection from Potential Abuse

The most obvious objection to a system of constant State surveillance is that it would be open to abuse. Jay Stanley and Barry Steinhardt, for example, note that the FBI's "Carnivore" program, which is "supposed to be used to tap into the email traffic of a particular individual," will filter through "all the traffic on the Internet Service Provider to which it has been attached"—and

they note, "the only thing that is keeping the government from trolling through all this traffic are software instructions written by the government itself."[20] Similarly, they note that under the PATRIOT Act the FBI need not show any probable cause (or even reasonable suspicion of criminal activity) to gain access to "individuals' financial records, medical histories, Internet usage, travel patterns, or … other records"—a situation that is ripe for abuse.[21] It must be admitted that, in practice, a system of constant State surveillance *is* likely to be abused to some extent.[22] However, if one adopts a rights-based understanding of claim (ii) above (i.e., if one holds that such a system of State surveillance is permissible as it does not in itself violate persons' moral rights), the theoretical force of this objection can be rebutted. On such an understanding of claim (ii) one could first note that this abuse-based objection gets its force from the view that such abuse would violate persons' moral rights. The proponent of a rights-based understanding of claim (ii) would certainly agree with this underlying view and would join with its advocates in condemning such abuse. However, the rights theorist who was in favor of such a system of State surveillance would also note that the condemnation of the *abuse* of State surveillance is not to condemn State surveillance *itself*. To condemn the use of *x* for the purposes of *y*, where *y* violates persons' rights (e.g., to privacy or autonomy), is *not* also to condemn the use of *x* for the purposes of *z*, where *z* does *not* violate persons' rights. Thus, a rights theorist who was a proponent of State surveillance could argue that to offer the possibility of abuse as an objection to constant State surveillance is to confuse the moral status of different possible uses of such surveillance. This is not, of course, either to downplay the badness of any abuse that might take place or to avoid the question of what practical measures should be imposed to attempt to minimize it. However, it is important to note that to object to a system of surveillance on the grounds that it could be abused is not to object to the surveillance itself, unless one believes (and can show) that the abuse in question would necessarily occur.

If one adopts a consequentialist understanding of claim (ii), however, defending the use of constant State surveillance against this objection is more difficult. This is because the likelihood of such abuse together with the likelihood of such abuse causing harm must be weighed against the benefits (as outlined above) that such a system is likely to provide. And, given that such a system has not yet been implemented, such a weighing and balancing of its relative costs and benefits will be difficult to assess with certainty. As such, then, the moral legitimacy of such a system would depend upon the political context in which it was proposed. If there is good reason to believe that little harm will accrue from the abuse of such a system of surveillance and there is reason to believe that it would bring important benefits to the citizens of the State in which it is installed, the possibility of its abuse should not deter consequentialists from endorsing the above pro-surveillance argument. Thus, it seems that installing a system of State surveillance in a State that possessed little power, a State that was too weak to exercise its power, or a State that had strict limits placed upon its power *and* that adhered to them would be justified on a consequentialist understanding of claim (ii). However, installing a system of State surveillance in a State that possessed considerable power over its citizens, that was strong enough to exercise such power, or that either had little restraints imposed upon it in its dealing with its citizens, or that had such restraints but frequently failed to adhere to them would be unlikely to be justified on a consequentialist understanding of claim (ii).[23]

Compelling Testimony Is Morally Impermissible

Both the rights theorists and the consequentialists who support the introduction of a system of constant State surveillance can thus meet the objection that such a system might be abused—although it should be noted that the consequentialist's response is such that the imposition of such a system would not be justified in some States. What, then, of the second objection that those who support the introduction of such a system of State surveillance are faced with: that it is not morally permissible for the agents of the State

(e.g., judges) to compel witness testimony or authorize information-gathering surveillance with probable cause? This second objection is less plausible than the first. As was noted in the initial clarificatory section of this essay, the above pro-surveillance argument is a *conditional* argument, such that *if* it is ever morally permissible for an agent of the State (e.g., a judge, or a prosecutor) to secure information from a witness about past events, *then* it is morally permissible for her to use State surveillance devices to secure the same information. For the antecedent clause of this argument to be true all that needs to be true is that *in at least one case* it would be morally permissible for a judge to secure information from a witness. And this latter claim is very plausible indeed, for both rights theorists and consequentialists alike. To see this, assume that Lee Harvey Oswald really was innocent of the assassination of John F. Kennedy, and, to prove his innocence, needed to secure the photograph that was taken of him standing in the doorway to the Texas School Book Depository at the time of the shooting. Assume also that the owner of this photograph disliked Oswald, and wanted to see him wrongly convicted. Here, it would clearly be morally permissible for the judge in Oswald's case to require that the owner of the photograph produce it in evidence. This claim becomes even stronger if the production of this photograph would cost its owner nothing but the frustration of his desire to see Oswald wrongly convicted; if, for example, the judge sent a messenger over to the house of the owner of the photograph to collect a copy of it, that he paid for this, and so on. Of course, extreme libertarians who oppose *any* form of interference by the State in the lives of its citizens would still hold that even this minimal interference in the life of the owner of the photograph is morally impermissible, and that under no circumstances should private individuals be required to give up any information, or any evidentiary items, that they might happen to possess.[24] It is not easy to provide a definitive rebuttal to this extreme libertarian position.[25] However, the advantages that this strict libertarian position possesses in terms of its ease of defense consequentialists alike), greatly outweighed by its highly counterintuitive results. (Or, at least, results that are counterintuitive to everyone *but* an extreme libertarian!) It is simple to construct examples in which the ill effects that result from a person's refusal to produce an item of information in evidence are vast, and where the costs to him of supplying this information are tiny. Imagine, for example, a case in which ten men will be executed for a crime that they did not commit unless a man who for idiosyncratic reasons of his own dislikes discussing what clothes he wore in the past discloses to their judge which clothes he wore on the day of the crime. In this case the clothes that the uncooperative witness wore are, for some reason, crucial to establishing whether it is the accused, or another group of ten men, who committed the crime in question. Clearly, in this case it would be counterintuitive to claim that the judge was not justified in requiring this man to reveal what clothes he wore on the day in question. As examples such as this multiply against the extreme libertarian's position, with the ill effects of the failure to produce the information increasing, and the costs involved in its provision decreasing, the claim that it is not morally permissible for a judge to impose the costs in question on the person withholding the information becomes increasingly implausible. And, once one accepts that for *any* case outlined in an example such as this, it would be morally permissible for a judge to secure the information in question, then one will have granted the truth of the antecedent claim on which this conditional argument in favor of surveillance rests. That is, once one grants that for any given case that it would be permissible for a judge to secure the withheld information, then one must also grant that it would be permissible to use State surveillance to secure the same information. And, since the State cannot tell where or when the actions or events that this information pertains to would occur, one would then also have to grant that the State would be justified in gathering information about *all* actions and events that occur, with, of course, the proviso that it can only *access* that which it is morally permissible for it to access. Unless one wishes to endorse the

that under no circumstances at all can a judge secure information from a witness that is relevant to a case in his court, then, one should reject this second objection to the above pro-surveillance argument.

SURVEILLANCE, PRIVACY, AND AUTONOMY

The two most obvious and immediate objections to the above pro-surveillance argument can thus be met. However, it still faces two more. The first of these is based on the claim that such a surveillance system would violate the privacy of those subjected to it. The second is based on the claim that it would illegitimately compromise their autonomy.[26]

Both of these objections have been leveled by Stanley and Steinhardt against the increase in surveillance in the United States under the administration of George W. Bush. According to Stanley and Steinhardt, both "Privacy and liberty in the United States are at risk," as it is "at risk of turning into a Surveillance Society."[27] As was noted above in discussing the consequentialist understanding of claim (ii) in response to the objection that a system of State surveillance of the type outlined above could be abused, the worries that Stanley and Steinhardt express would be legitimate ones in certain political contexts—including that which they are explicitly addressing. However, were the system of surveillance that is outlined above to be installed in a State that would be unlikely to abuse it, neither of their objections would be sound.

The objection that constant State surveillance would violate the privacy of those subject to it rests on a failure to acknowledge the legitimate limits of the proposed surveillance. On the system of surveillance argued for above the State would only be permitted to access the information that its surveillance devices gather *when it is morally permissible for it to do so*. As such, the State's use of such a surveillance system would not enable it justly to access any more information than it is morally permitted to access already. This point is not reiterated to defend the system of State surveillance advocated above on the grounds that it would not violate the privacy of its citizens any more than this is violated already. Rather, it has been reiterated to provide the basis for a *stronger* defense of this system of State surveillance: *that it would not violate the citizens' privacy at all.*

To develop this strong defense of the system of State surveillance advocated above it must be recognized that privacy is both a relative notion and a normative notion. A certain item of information is private *relative to* a certain person, if that person cannot *legitimately* (a normative concept) require that that item of information be disclosed to her.[28] For example, my use of the checking account that I share with my wife is not private with respect to her, for she can legitimately require that I disclose my use of it to her. However, it is private with respect to my colleagues, who cannot legitimately require that I disclose my use of it to them. Similarly, the examination score of a student in my class is not private with respect to the student, since she can legitimately require its disclosure, but it is private with respect to my wife, who cannot. Given, then, that privacy is both a *relative* notion and a *normative* notion, and since, according to the pro-surveillance argument above, the agents of the State can only access those items of information that they are morally permitted to access, it is clear that the above pro-surveillance argument does not justify the violation of the privacy of the State's citizens. According to that argument, the only time at which the State could permissibly access information about actions or events would be when it was morally permitted to do so; that is, when it could *legitimately* require that this information be disclosed to it. This being so, under the system of State surveillance argued for above, the State could only access information that is *not* private relative to it. Thus, in accessing such information the State would not violate the privacy of its citizens.[29]

Just as constant and universal State surveillance of the sort argued for above would not violate the privacy of the persons subject to it, nor would it compromise their autonomy—although it is easy to see why one might think that it would. In *1984* Oceania's citizens adjusted their behavior when they were in the presence of telescreens to

conform to the way that they believed that Big Brother's Party wanted them to behave.[30] Were they not to do this, and were they to be observed by agents of Big Brother during their refusal to conform, they risked severe punishment. To avoid the penalties that the Party would impose on them if they failed to conform to its expectations, then, Oceania's citizens ceded a degree of control over their actions to the Party. That is, when they acted out of fear that Party members were observing them, and so acted to conform to the Party's view of how they should behave, the citizens of Oceania satisfied their individual first-order desires to "Perform the actions that the Party wants me to perform." To the extent that they thus ceded control to the Party, then, Oceania's citizens were other-directed, and so heteronomous, rather than being self-directed, and so autonomous.[31] Thus, in *1984*, Big Bother's citizens suffered from compromised autonomy as a result of being subjected to (or potentially subjected to) State surveillance.[32]

Fortunately, however, the citizens of a State that utilized the type of constant surveillance that would be justified on the above pro-surveillance argument would not similarly suffer from compromised autonomy. As noted earlier such a surveillance system would not utilize the type of in-person surveillance that was used by Big Brother's Party in *1984*. Instead, it would merely record the actions of the citizens and the events that they participated in. Since this is so, most of the citizens subject to this form of State surveillance would not believe that they needed to alter their behavior to conform to the State's view of how they should act (were the State even to hold such a view) for they would recognize that the State would not actually be watching them. Moreover, if they were law-abiding then they would also recognize that the chances of the State having just cause to access information about their actions would be very slim indeed. Free from the pressure to conform of the sort that was imposed upon the citizens of Oceania, then, most citizens whose State subjected them to constant and universal surveillance of the type outlined above would *not* cede control over their acts to the State. The citizens of such a State would thus not become heteronomous with respect to their behavior, as did the citizens of Oceania. Instead, they would retain their autonomy with respect to it. Moreover, the claim that most citizens placed under constant State surveillance would not suffer from any diminution in their autonomy is not merely a speculative one. Persons subjected to surveillance are unlikely to alter their behavior once they become used to being "on tape" (unless they previously performed acts that they believed that they should not have been performing) when they realize that there is little chance that the information that the surveillance devices record would ever be required, and so there is little chance of their actions ever being observed.[33]

Of course, not all the citizens of a State that utilized the type of surveillance advocated above would retain full autonomy with respect to all of their actions. The exceptions would be those who are criminally inclined, and who, as a result of being placed under State surveillance, would alter their behavior by refraining from committing crimes. Such persons would suffer from compromised autonomy with respect to their deliberate omission of criminal activity, insofar as they alter their behavior solely to avoid incurring criminal penalties. Yet, given that the actions that these persons refrain from performing are criminal ones, the diminution in autonomy that these persons would experience as a result of being placed under State surveillance is one that it is morally legitimate to inflict upon them.[34]

CONCLUSION

It is now time to take stock. Although it is often claimed that the recent proliferation of surveillance technology is turning the West into an "Orwellian nightmare" it was argued in this paper that, rather than condemning the prospect of constant State surveillance, we should instead welcome it—provided that we lived in a State where its abuse would be unlikely.[35] Such surveillance would not, however, involve any expansion of State power into the lives of its citizens. If, under certain circumstances, it is morally permissible for judges to secure information relevant to their cases from witnesses, then, under the same

circumstances, it should also be morally permissible for them to secure this information through the use of surveillance devices. Constant State surveillance is thus no different in principle than the current system of subpoenaing witnesses. Moreover, the State's use of such surveillance would be morally *preferable* to the subpoenaing of witnesses. It would, for example, avoid the need to impose costs on those who would otherwise be called as witnesses, and it would result in the provision of more accurate information. There are, of course, practical concerns that must be addressed before any such State surveillance system is put in place. But since, as has been argued in this paper, there is no *principled* reason to oppose such a system, but there is reason to endorse it, once worries about the possibility of such a system surveillance being abused are laid to rest the road would be clear for its introduction. Big Brother would then indeed be watching over us. But, unless one has criminal tendencies, this should be a cause for relief, rather than concern.[36]

NOTES

1. George Orwell, 1984 (London: Seeker and Warburg, 1987), 119.
2. Ibid., 123.
3. I thank an anonymous referee for pressing me on this point.
4. Note, however, that owing to the possibility that a constant and universal system of State surveillance will be subject to abuse it will be easier to justify such a system if one adopts a rights-based approach to ethics than if one adopts a consequentialist approach. This is discussed more fully below.
5. At least, if one accept a plausible version of consequentialism, such that human well-being is the good to be maximized.
6. A further clarification is also in order. Since claim (ii) can be accepted by both rights theorists and consequentialists, the fact that the arguments for the view that a system of State surveillance is morally preferable to the absence of such a system are based on the good consequences that such a system would have does not preclude rights theorists from accepting these arguments. A rights theorist could endorse such consequentialist reasoning, provided that such a system of State surveillance would respect those moral rights that she believes that persons possess.
7. That Anglo-American common law recognizes these exemptions to personal privacy was noted by Warren and Brandeis in their seminal article "The right to privacy [The implicit made explicit]," in Ferdinand D. Schoeman, ed., *Philosophical Dimensions of Privacy: An Anthology* (Cambridge, Cambridge University Press, 1984), 78, 88–89. Warren's and Brandeis's approval of, and lack of defense of, such exemptions indicates both that they found them to be morally permissible, and also that they assumed (as it is also assumed in the arguments in this paper) that this view would be so widely accepted as to need no defense.
8. Although the claim that, under certain circumstances, judges are morally permitted to compel witnesses to testify and to permit law enforcement agencies to monitor suspects is innocuous, the question of under what circumstances such compulsion and surveillance is morally permissible is a contentious one. As will be noted below, however, the arguments in this paper rest only on the innocuous view *that in some circumstances* such compulsion and surveillance is morally permissible, and not on any substantive claims concerning *when* such information gathering is morally permissible.
9. See, for example, Robert J. Araujo, "International Tribunals and Rules of Evidence: The Case for Respecting and Preserving the 'Priest-Penitent' Privilege Under International Law," *American University International Law Review* 15 (2000): 639–666.
10. See, for example, Jean V. McHale, "Medical Confidentiality and Legal Privilege," *Journal of Applied Philosophy* 11, no. 2(1994): 241–242.
11. See, for example, Paul R. Rice, "Attorney-Client Privilege: The Eroding Concept of Confidentiality Should Be Abolished," *Duke Law Journal* 47 (1998): 853–898, and John R. Przypyszny, "Public Assault on the Attorney-Client Privilege: Ramifications of Bakes v. Doe," *Georgetown Journal of Legal Ethics* 3 (1989): 351.
12. See, for example, Alan Meisel, "Confidentiality and Rape Counseling—Privacy vs. the Right to a Fair Trial," *Hastings Center Report* 11 (1981): 5–7.
13. See, for example, both Mark R. Wicclair, "A Shield Privilege for Reporters vs. the Administration of Justice and the Right to a Fair Trial: Is There a Conflict?," *Business and Professional Ethics Journal* 4 (1985): 1–14, and Richard P. Cunningham, "Commentary on Wicclair's: A Shield Privilege for Reporters vs. the Administration of Justice

and the Right to a Fair Trial," *Business and Professional Ethics Journal* 4 (1985): 15–17.

14. For a general discussion of this issue, see Bernard Baumrin, "Is There a Freedom Not to Speak?" *Metaphilosophy* 6 (1975): 25–34.
15. Thus, this argument should *not* be taken as an endorsement of the extended surveillance powers granted to law enforcement agencies by this Act.
16. The use of such technology was noted by Andrew J. Charlesworth, "Privacy, Personal Information, and Employment," *Surveillance and Society* 1, no. 2 (2003): 218.
17. Jeremy Bentham's Panopticon was originally designed to be a humane prison, in which the prisoners were reformed through subjecting them both to isolation and to the possibility of covert surveillance. Such conditions, Bentham thought, would lead the prisoners to modify their own behavior to conform to that which they believed their (possible) observers would endorse. The Panopticon was designed as a ring of cells, in the center of which was an observation tower equipped with special shutters, so that the guards could see out, but the prisoners could not see in—and so could not tell if they were being watched or not. See Jeremy Bentham, *Panopticon, or, The Inspection House*, originally published 1791, and Janet Semple, *Bentham's Prison: A Study of the Panopticon Penitentiary* (Oxford Clarendon Press, 1993).
18. That overt surveillance deters persons from committing crimes is well documented. J. Ditton and E. Short, for example, have shown that the use of CCTV surveillance in Airdrie, Scotland, resulted in a significant reduction in crime in the immediate area, with no displacement of criminal activity to the adjoining areas without CCTV. They also noted that after the introduction of CCTV there was a reduction in the number of large-scale affrays. "Evaluating Scotland's first town center CCTV scheme," in C. Norris, J. Moran, and G. Armstrong, eds., *Surveillance, Closed Circuit Television and Social Control* (Aldershot, UK: Ashgate, 1998), 155–174.
19. In a survey on the effect that the presence of CCTV had on police behavior that was conducted by Benjamin J. Goold between June 1997 and March 2000 in six towns in the south of England over two-thirds of the fifty police officers interviewed said that "the introduction of cameras patrol" to avoid being prosecuted for such breaches of duty as unlawful arrest or assault. See Benjamin J. Goold, "Public Area Surveillance and Police Work: The Impact of CCTV on Police Behavior and Autonomy," *Surveillance & Society* 1, no. 2 (2002) 192, 194.
20. Jay Stanley and Barry Steinhardt, *Bigger Monster, Weaker Chains: The Growth of an American Surveillance Society* (ACLU Technology and Liberty Program, 2003), 8.
21. Ibid., 9.
22. Benjamin J. Goold claims that 1% of the targets tracked by CCTV operators in Britain are for voyeuristic purposes. See his *CCTV and Policing Public Area Surveillance and Police Practices in Britain* (Oxford: Oxford University Press, 2004), 144. It must be stressed that the use of CCTV that Goold examined was not subject to the penalties that are proposed for misuse in this paper.
23. As such, then, it is highly unlikely that State surveillance of any but the most minimal kind would be justified in the contemporary United States, owing to the lack of judicial oversight of State action that Stanley and Steinhardt note.
24. Few rights theorists ascribe to this extreme libertarian position. Instead, the rights that they ascribe to persons such as privacy, or private property, are frequently conditional on such persons not being situated in "emergency situations" such as this.
25. One might, of course, challenge the extreme libertarian to provide a defensible account of the strong property rights that she must invoke as a basis for her argument. But this would not provide a definitive rebuttal of the extreme libertarian position so much as it would merely shift the burden of proof to its proponents.
26. These two objections might also be expressed in terms of moral rights, namely, that such surveillance would violate the citizens' right to privacy, or their right to autonomy. If the following arguments are sound, then, they also show that such rights-based objections are mistaken. See also notes 31 and 36.
27. Stanley and Steinhardt, *Bigger Monster, Weaker Chains*, p. 1. Although Stanley and Steinhardt use the term "liberty" rather than "autonomy," they make it clear in their essay that they are concerned with the erosion of control that persons have over how they live their lives that could be

and to the increase in the possibility that persons might be subject to an increase in State coercion if they are placed under surveillance. As such, then, it would not be inappropriate to hold that their concern with liberty is also a concern with autonomy.

28. The account of privacy that is offered here is an account of what it is for something (e.g., an item of information) to *be* private, rather than an account of what it is for a person to be *in a condition* of privacy. (As Adam Moore puts it, the concern of the account of privacy outlined here "is what should be considered a 'private affair'—something that is no one else's business." "Intangible Property: Privacy, Power, and Information Control," *American Philosophical Quarterly*, Vol. 35, no. 4 [1998], 372.) The focus of this account of privacy is thus different from accounts of privacy such as that developed by William Parent or Ferdinand Schoeman, for whom, respectively, "Privacy is the condition of not having undocumented personal knowledge about one possessed by others," and "A person has privacy to the extent that others have limited access to information about him … the intimacies of his life, or … his thoughts or his body." W. A. Parent, "Privacy, Morality, and the Law," in D. Johnson and J. Snapper, *Ethical Issues in the Use of Computers* (Wadsworth, 1985), 203 (reprinted from *Philosophy and Public Affairs* [1983]: 269–288), and Ferdinand Schoeman, "Privacy Philosophical Dimensions of the Literature," in Schoeman, ed., *Philosophical Dimensions of Privacy*, 3. This difference in focus is important, for it explains why descriptive (rather than normative) accounts of privacy (such as Parent's and Schoeman's) do not compete with the normatively laden account that is offered here, and so there is no need to argue that this account is preferable to (e.g.) Parent's.

29. Thus, if persons did possess a moral right to privacy this pro-surveillance argument would not justify its violation. And this is true whether this right to privacy was purely a principle-based right, or one that was based on practical considerations (i.e., that such a right is needed to protect some more fundamental human good, such as autonomy or well-being).

30. It should be noted that, unlike the type of surveillance argued for here, that which the Party subjected the citizens of Oceania to frequently *did* violate their privacy.

31. For a discussion of the relationships that hold between privacy, autonomy, and the ceding of control, see Taylor, "Autonomy and Privacy," 460–473.

32. "There was of course no way of knowing whether you were being watched at any given moment.… You had to live—did live, from habit that became instinct—in the assumption that every sound you made was overheard, and, except in darkness, every movement scrutinized," Orwell, 1984, 4–5.

33. David Lyon offers evidence that shows that persons do not change their behavior when they are under surveillance once they realize that it is unlikely that their images will be accessed. See his discussion of the use of surveillance cameras in retail stores in *Surveillance Society: Monitoring Everyday Life* (Milton Keynes, UK: Open University Press), 51, 62.

34. Thus, if persons did possess a moral right to autonomy, this pro-surveillance argument would not justify its violation.

35. The claim that the proliferation of surveillance devices will lead to an "Orwellian nightmare" occurs frequently in the philosophical literature on this subject. See, for example, Adam Moore, "Employee Monitoring and Computer Technology: Evaluative Surveillance V. Privacy," *Business Ethics Quarterly*, Vol. 10, no. 3 (2000), 698. 36.

36. I thank an anonymous referee for *Public Affairs Quarterly* for his or her exceptionally generous and helpful comments on an earlier version of this paper, and the Editor of that journal for encouraging me to develop my arguments. I also thank Mort Winston for encouraging me to refine this paper further, and to acknowledge the concerns that would legitimately arise from the increase in State surveillance in the present day United States.

2.2 ARTIFICIAL INTELLIGENCE AND ROBOTICS

2.2.1 BIRTH OF A ROBOT

ABIGAIL TUCKER

Early descriptions of life-like mechanical devices go back to ancient Greece with more modern ones including automatons which played musical instruments or even chess. The later device, named the "chess Turk," was invented in Germany toward the end of the eighteenth century, but was discredited after many years when it was disclosed that a small person was hidden inside directing its operation. Mary Shelley's famous work of science fiction, *Frankenstein or the Modern Prometheus*, first published anonymously in 1818, described an experiment to create life by assembling parts from cadavers and then bringing the reconstructed monster to life using a bolt of lightning to jumpstart its heart.

In 2000 Honda introduced the humanoid robot ASIMO (Advanced Step in Innovative Mobility) which was developed in order to assist people with mobility impairments. Making robots that perform different physical functions, such as welding, is relatively simple compared to creating a machine that is able to interact with a person in a human-like manner. A *social robot* should be able to acquire knowledge and learn on its own and then use that information in a situation which it has never encountered previously.

In this selection, Abigail Tucker describes the work of a team of robotidists, engineers, and developmental cognitive psychologists who are working to create a social robot that approximates the sensory and motor capabilities of a human infant. Achieving even this level of complexity in a robot is proving to be exceedingly difficult, but by pursuing the ambitious goals of this project we are learning a lot about both robots and ourselves.

FOCUS QUESTIONS

1. What are the key differences between the human and machine "learning" described in this article?
2. Robots and automated machines are usually controlled by some type of programming, ranging from the simple to the complex. How does the author describe the differences between the "learning" behaviors of an advanced social robot and those of a human being? How might each make decisions about how and when to act?
3. Why might social robots have appeal in some societies and how could they be useful, if they ever reach a sufficiently high level of development?
4. What are some of the obstacles that must be overcome before humans are likely to accept social robots in their lives? If that ever happens, what ethical issues might arise and need to be resolved? Discuss.

Source: Abigail Tucker, "Birth of a Robot," *Smithsonian* magazine, Volume 40, Issue 4, Jul 2009, pp. 56-65. Copyright 2012 Smithsonian Institution. Reprinted with permission from Smithsonian Enterprises. All rights reserved. Reproduction in any medium is strictly prohibited without permission from Smithsonian Institution. Such permission may be requested from Smithsonian Enterprises.

KEYWORDS

artificial intelligence, bootstrapping, cognitive scientist, conversational skill, conversational software, intelligent system, social fluency, social robot

Einstein the robot has enchanting eyes, the color of honey in sunlight. They are fringed with drugstore-variety false eyelashes and framed by matted gray brows made from real human hair. "What is that, makeup?" a visiting engineer asks, and, indeed, on closer examination I can see black eyeliner smeared beneath Einstein's lower lids, a la David Bowie in 1971. The machine's gaze is expressive—soulful, almost.

David Hanson, Einstein's creator, is visiting from Texas to help scientists here at the University of California at San Diego (UCSD) prepare the robot for an upcoming conference. Hanson switches the robot on—really just a head and neck—and runs it through some of its dozens of expressions. Its lips purse. Its brow furrows. Its eyes widen as though in horror, then scrunch mirthfully as it flashes a grin. The 27 motors in the face make a wretched grinding sound, and when the mouth opens I see a tangle of wires where the prodigious brain should be. Einstein's white wig is missing and the skin of its neck hangs in flaps, because its shoulders, made of plastic, got shattered in shipping.

Still, the effect is so lifelike that even jaded graduate students have stopped by to stroke the robot's wrinkled cheek, which is encased in a substance called Frubber—"face rubber," Hanson explains—that feels buttery soft and cold as a cadaver. Engineers working at night have been so unnerved by Einstein's presence they asked that it be covered up.

That delights Javier Movellan, leader of the university group that recently spent $70,000 in research grants to buy the robotic head from Hanson. Einstein will eventually be able to track people's movements with cameras that twinkle behind its pupils and judge their gender, age (well, whether they're over or under 40), if they wear glasses or have facial hair and when they're smiling. The robot is also equipped with conversational software and can hold forth while staring deep into a person's eyes. "Do you know any jokes?" Hanson asks the machine.

Einstein: Why did the chicken cross the road?
Hanson: Tell me about artificial intelligence.
Einstein: I feel real, not artificial.
Hanson: I want to hear more about the future of robotics.
Einstein: Robots will become ever more like animals and people. Robots will continue to get more amazing and cool.

Einstein is the product of a remarkable collaboration. Hanson, a robot designer and the founder of the Dallas-based firm Hanson Robotics, has used classical sculpting techniques to animate robotic likenesses of Philip K. Dick, author of Do Androids Dream of Electric Sheep? (the basis of the apocalyptic movie Blade Runner), his own wife (he had to use a male skull model, "which masculinized her a bit") and more than a dozen other people. Movellan, a psychologist and software pioneer who runs UCSD's Machine Perception Laboratory, develops technology that approximates human senses. Einstein is, at present, a research tool to explore how a machine can perceive and react to human facial expressions; that capacity could later have many practical applications in entertainment and education, alerting the robot teachers of the future, say, that their human pupils are daydreaming.

For the most part, though, the intelligence I perceived in Einstein—its intense eye contact, its articulate soliloquies—was an illusion. Its answers to questions were canned and its interpretive powers were extremely limited. In short, Einstein is no Einstein. Overall, robots can do amazing things—play the violin, dismantle bombs, fire missiles, diagnose diseases, tend tomato plants, dance—but they sorely lack the basics. They recite jokes but don't get them. They can't summarize a movie. They can't tie their shoelaces. Because of such shortcomings, whenever we encounter them in the flesh, or Frubber, as it were, they are bound to disappoint.

Rodney Brooks, an M.I.T. computer scientist who masterminded a series of robotics innovations

in the 1990s, said recently that for a robot to have truly humanlike intelligence, it would need the object-recognition skills of a 2-year-old child, the language capabilities of a 4-year-old, the manual dexterity of a 6-year-old and the social understanding of an 8-year-old. Experts say they are far from reaching those goals. In fact, the problems that now confound robot programmers are puzzles that human infants often solve before their first birthday. How to reach for an object. How to identify a few individuals. How to tell a stuffed animal from a bottle of formula. In babies, these skills are not preprogrammed, as were the perceptual and conversational tricks Einstein showed me, but rather are cultivated through interactions with people and the environment.

But what if a robot could develop that way? What if a machine could learn like a child, as it goes along? Armed with a nearly $3 million National Science Foundation grant, Movellan is now tackling that very question, leading a team of cognitive scientists, engineers, developmental psychologists and roboticists from UCSD and beyond. Their experiment—called Project One, because it focuses on the first year of development—is a wildly ambitious effort to crack the secrets of human intelligence. It involves, their grant proposal says, "an integrated system ... whose sensors and actuators approximate the levels of complexity of human infants."

IN OTHER WORDS, A BABY ROBOT

The word "ROBOT" hit the world stage in 1921, in the Czech science fiction writer Karel Capek's play Rossum's Universal Robots, about a factory that creates artificial people. The root is the Czech robota, for serf labor or drudgery. Broadly understood, a robot is a machine that can be programmed to interact with its surroundings, usually to do physical work.

We may associate robots with artificial intelligence, which uses powerful computers to solve big problems, but robots are not usually designed with such lofty aspirations; we might dream of Rosie, the chatty robot housekeeper on "The Jetsons," but for now we're stuck with Roomba, the disk-shaped, commercially available autonomous vacuum cleaner. The first industrial robot, called Unimate, was installed in a General Motors factory in 1961 to stack hot pieces of metal from a die-casting machine. Today, most of the world's estimated 6.5 million robots perform similarly mundane industrial jobs or domestic chores, though 2 million plug away at more whimsical tasks, like mixing cocktails. "Does [the robot] prepare the drink with style or dramatic flair?" ask the judging guidelines for the annual RoboGames bartending competition, held in San Francisco this summer. "Can it prepare more than a martini?"

Now imagine a bartender robot that could waggle its eyebrows sympathetically as you pour out the story of your messy divorce. Increasingly, the labor we want from robots involves social fluency, conversational skill and a convincing humanlike presence. Such machines, known as social robots, are on the horizon in health care, law enforcement, child care and entertainment, where they might work in concert with other robots and human supervisors. Someday, they might assist the blind; they've already coached dieters in an experiment in Boston. The South Korean government has said it aims to have a robot working in every home by 2020.

Part of the new emphasis on social functioning reflects the changing economies of the richest nations, where manufacturing has declined and service industries are increasingly important. Not coincidentally, societies with low birthrates and long life expectancies, notably Japan, are pushing hardest for social robots, which may be called upon to stand in for young people and perform a wide variety of jobs, including caring for and comforting the old.

Some scientists working on social robots, like Movellan and his team, borrow readily from developmental psychology A machine might acquire skills as a human child does by starting with a few basic tasks and gradually constructing a more sophisticated competence—"bootstrapping," in scientific parlance. In contrast to preprogramming a robot to perform a fixed set of actions, endowing a robot computer with the capacity to acquire skills gradually in response to the environment might produce smarter, more human robots.

"If you want to build an intelligent system, you have to build a system that becomes intelligent," says Giulio Sandini, a bioengineer specializing in socialrobots at the Italian Institute of Technology in Genoa. "Intelligence is not only what you know but how you learn more from what you know. Intelligence is acquiring information, a dynamic process."

"THIS IS THE BRAINS!" Movellan shouted over the din of cyclone-strength air conditioners. He was pointing at a stack of computers about ten feet tall and six feet deep, sporting dozens of blinking blue lights and a single ominous orange one. Because the Project One robot's metal cranium will not be able to hold all the information-processing hardware that it will need, the robot will be connected by fiber-optic cables to these computers in the basement of a building on the UCSD campus in La Jolla. The room, filled with towering computers that would overheat if the space weren't kept as cold as a meat locker, looks like something out of *2001: A Space Odyssey*.

As Einstein could tell you, Movellan is over 40, bespectacled and beardless. But Einstein has no way of knowing that Movellan has bright eyes and a bulky chin, is the adoring father of an 11-year-old daughter and an 8-year-old son and speaks English with an accent reflecting his Spanish origins.

Movellan grew up amid the wheat fields of Palencia, Spain, the son of an apple farmer. Surrounded by animals, he spent endless hours wondering how their minds worked. "I asked my mother, 'Do dogs think? Do rats think?'" he says. "I was fascinated by things that think but have no language."

He also acquired a farm boy's knack for working with his hands; he recalls that his grandmother scolded him for dissecting her kitchen appliances. Enamored of the nameless robot from the 1960s television show "Lost in Space," he built his first humanoid when he was about 10, using "food cans, light bulbs and a tape recorder," he says. The robot, which had a money slot, would demand the equivalent of $100. As Movellan anticipated, people usually forked over much less. "That's not $100!" the robot's prerecorded voice would bellow. Ever the mischievous tinkerer, he drew fire 30 years later from his La Jolla homeowners association for welding robots in his garage.

He got his PhD in developmental psychology at the University of California at Berkeley in 1989 and moved on to Carnegie Mellon University, in Pittsburgh, to conduct artificial intelligence research. "The people I knew were not really working on social robots," he says. "They were working on vehicles to go to Mars. It didn't really appeal to me. I always felt robotics and psychology should be more together than they originally were." It was after he went to UCSD in 1992 that he began working on replicating human senses in machines.

A turning point came in 2002, when he was living with his family in Kyoto, Japan, and working in a government robotics lab to program a long-armed social robot named Robovie. He hadn't yet had much exposure to the latest social robots and initially found them somewhat annoying. "They would say things like, 'I'm lonely, please hug me,'" Movellan recalls. But the Japanese scientists warned him that Robovie was special. "They would say, 'you'll feel something.' Well, I dismissed it—until I felt something. The robot kept talking to me. The robot looked up at me and, for a moment, I swear this robot was alive."

Then Robovie enfolded him in a hug and suddenly—"magic," says Movellan. "This is something I was unprepared for from a scientific point of view. This intense feeling caught me off guard. I thought, Why is my brain put together so that this machine got me? Magic is when the robot is looking at things and you reflexively want to look in the same direction as the robot. When the robot is looking at you instead of through you. It's a feeling that comes and goes. We don't know how to make it happen. But we have all the ingredients to make it happen."

Eager to understand this curious reaction, Movellan introduced Robovie to his 2-year-old son's preschool class. But there the robot cast a different spell. "It was a big disaster," Movellan remembers, shaking his head. "It was horrible. It was one of the worst days of my life." The toddlers were terrified of Robovie, who was

about the size of a 12-year-old. They ran away from it screaming.

That night, his son had a nightmare. Movellan heard him muttering Japanese in his sleep: "Kowai, kowai." Scary, scary

Back in California, Movellan assembled, in consultation with his son, a kid-friendly robot named RUBI that was more appropriate for visits to toddler classrooms. It was an early version of the smiling little machine that stands sentinel in the laboratory today, wearing a jaunty orange Harley-Davidson bandanna and New Balance sneakers, its head swiveling in an inquisitive manner. It has coasters for eyes and a metal briefcase for a body that snaps open to reveal a bellyful of motors and wires.

"We have learned a lot from this little baby," Movellan said, giving the robot an affectionate pat on its square cheek.

For the past several years he has embedded RUBI at a university preschool to study how the toddlers respond. Various versions of RUBI (some of them autonomous and others puppeteered by humans) have performed different tasks. One taught vocabulary words. Another accompanied the class on nature walks. (That model was not a success; with its big wheels and powerful motors, RUBI swelled to an intimidating 300 pounds. The kids were wary, and Movellan was, too.)

The project has had its triumphs—the kids improved their vocabularies playing word games displayed on RUBI's stomach screen—but there have been setbacks. The children destroyed a fancy robotic arm that had taken Movellan and his students three months to build, and RUBI's face detector consistently confused Thomas the Tank Engine with a person. Programming in incremental fixes for these problems proved frustrating for the scientists. "To survive in a social environment, to sustain interaction with people, you can't possibly have everything preprogrammed," Movellan says.

Those magic moments when a machine seems to share in our reality can sometimes be achieved by brute computing force. For instance, Einstein's smile-detection system, a version of which is thousands of photographs of faces that had been marked "smiling" or "not smiling." After cataloging those images and discerning a pattern, Einstein's computer can "see" whether you are smiling, and to what degree. When its voice software is cued to compliment your pretty smile or ask why you look sad, you might feel a spark of unexpected emotion.

But this laborious analysis of spoon-fed data—called "supervised learning"—is nothing like the way human babies actually learn. "When you're little nobody points out ten thousand faces and says 'This is happy, this is not happy, this is the left eye, this is the right eye,'" said Nicholas Butko, a PhD student in Movellan's group. (As an undergraduate, he was sentenced to labeling a seemingly infinite number of photographs for a computer face-recognition system.) Yet babies are somehow able to glean what a human face is, what a smile signifies and that a certain pattern of light and shadow is Mommy.

To show me how the Project One robot might learn like an infant, Butko introduced me to Bev, actually BEV, as in Baby's Eye View I had seen Bev slumped on a shelf above Butko's desk without realizing that the Toys 'R' Us-bought baby doll was a primitive robot. Then I noticed the camera planted in the middle of Bev's forehead, like a third eye, and the microphone and speaker under its purple T-shirt, which read, "Have Fun."

In one experiment, the robot was programmed to monitor noise in a room that people periodically entered. They'd been taught to interact with the robot, which was tethered to a laptop. Every now and then, Bev emitted a babylike cry. Whenever someone made a sound in response, the robot's camera snapped a picture. The robot sometimes took a picture if it heard no sound in response to its cry, whether or not there was a person in the room. The robot processed those images and quickly discerned that some pictures—usually those taken when it heard a response—included objects (faces and bodies) not present in other pictures. Although the robot had previously been given no information about human beings (not even that such things existed),

someone was in the room. In a remarkably short time, Bev had "discovered" people.

A similar process of "unsupervised learning" is at the heart of Project One. But Project One's robot will be much more physically sophisticated than Bev—it will be able to move its limbs, train its cameras on "interesting" stimuli and receive readings from sensors throughout its body—which will enable it to borrow more behavior strategies from real infants, such as how to communicate with a caregiver. For example, Project One researchers plan to study human babies playing peekaboo and other games with their mothers in a lab. Millisecond by millisecond, the researchers will analyze the babies' movements and reactions. This data will be used to develop theories and eventually programs to engineer similar behaviors in the robot.

It's even harder than it sounds; playing peekaboo requires a relatively nuanced understanding of "others." "We know it's a hell of a problem," says Movellan. "This is the kind of intelligence we're absolutely baffled by. What's amazing is that infants effortlessly solve it." In children, such learning is mediated by the countless connections that brain cells, or neurons, form with one another. In the Project One robot and others, the software itself is formulated to mimic "neural networks" like those in the brain, and the theory is that the robot will be able to learn new things virtually on its own.

The robot baby will be able to touch, grab and shake objects, and the researchers hope that it will be able to "discover" as many as 100 different objects that infants might encounter, from toys to caregivers' hands, and figure out how to manipulate them. The subtleties are numerous; it will need to figure out that, say, a red rattle and a red bottle are different things and that a red rattle and a blue rattle are essentially the same. The researchers also want the robot to learn to crawl and ultimately walk.

Perhaps the team's grandest goal is to give the robot the capacity to signal for a caregiver to retrieve an object beyond its grasp. Movellan calls this the "Vygotsky reach," after developmental psychologist Lev Vygotsky, who identified the about a year old—as an intellectual breakthrough, a transition from simple sensory-motor intelligence to symbolic intelligence. If the scientists are successful, it will be the first spontaneous symbolic gesture by a robot. It will also be a curious role reversal—the robot commanding the human, instead of vice versa.

"That's a pretty important transition," says Jonathan Plucker, a cognitive scientist at Indiana University who studies human intelligence and creativity. Plucker had no prior knowledge of Project One and its goals, but he was fresh from watching the season finale of "Battlestar Galactica," which had left him leery of the quest to build intelligent robots. "My sense is that it wouldn't be hard to have a robot that reaches for certain types of objects," he says, "but it's a big leap to have a machine that realizes it wants to reach for something and uses another object, a caregiver, as a tool. That is a much, much more complex psychological process."

At present, the Project One robot is all brains. While the big computer hums in its air-conditioned cavern, the body is being designed and assembled in a factory in Japan. Construction is expected to take about nine months.

A prototype of the Project One robot body already exists, in the Osaka laboratory of Hiroshi Ishiguro, the legendary Japanese roboticist who, in addition to creating Robovie, fashioned a robotic double of himself, named Geminoid, as well as a mechanical twin of his 4-year-old daughter, which he calls "my daughter's copy." ("My daughter didn't like my daughter's copy," he told me over the phone. "Its movement was very like a zombie." Upon seeing it, his daughter—the original—cried.) Ishiguro's baby robot is called the Child-Robot with Biomimetic Body, or CB2 for short. If you search for "creepy robot baby" on YouTube, you can see clips of four-foot-tall CB2 in action. Its silicone skin has a grayish cast; its blank, black eyes dart back and forth. When first unveiled in 2007, it could do little more than writhe, albeit in a very babylike way, and make pathetic vowel sounds out of the

"It has this ghostly gaze," says Ian Fasel, a University of Arizona computer scientist and a former student of Movellan's who has worked on the Japanese project. "My friends who see it tell me to please put it out of its misery. It was often lying on the floor of the lab, flopping around. It gives you this feeling that it's struggling to be a real boy, but it doesn't know how."

When Movellan first saw CB2, last fall, as he was shopping around for a Project One body, he was dismayed by the lack of progress the Japanese scientists had made in getting it to move in a purposeful way "My first impression was that there was no way we would choose that robot," Movellan recalls. "Maybe this robot is impossible to control. If you were God himself, could you control it?"

Still, he couldn't deny that the CB2 was an exquisite piece of engineering. There have been other explicitly childlike robots over the years—creations such as Babybot and Infanoid—but none approach CB2's level of realism. Its skin is packed with sensors to collect data. Its metal skeleton and piston-driven muscles are limber, like a person's, not stiff like most robots', and highly interconnected: if an arm moves, motors in the torso and elsewhere respond. In the end, Movellan chose CB2.

The body's humanness would help the scientists develop more brainlike software, Movellan decided. "We could have chosen a robot that could already do a lot of the things we want it to do—use a standard robotic arm, for instance," Movellan says. "Yet we felt it was a good experiment in learning to control a more biologically inspired body that approximates how muscles work. Starting with an arm more like a real arm is going to teach us more."

The Project One team has requested tweaks in CB2's design, to build in more powerful muscles that Movellan hopes will give it the strength to walk on its own, which the Japanese scientists—who are busy developing a new model of their own—now realize the first CB2 will never do. Movellan is also doing away with the skin suit, which sometimes provides muddled readings, opting instead for a Terminator-like metal skeleton encased in clear plastic. ("You can always put clothes on," Movellan reasons.) He had hoped to make the robot small enough to cradle, but the Japanese designers told him that is currently impossible. The baby will arrive standing about three feet tall and weighing 150 pounds.

What a socialrobot's face should look like is a critical, and surprisingly difficult, decision. CB2's face is intended to be androgynous and abstract, but somehow it has tumbled into what robotics experts term the "uncanny valley," where a machine looks just human enough to be unsettling. The iCub, another precocious child-inspired robot being built by a pan-European team, looks more appealing, with cartoonish wide eyes and an endearing expression. "We told the designers to make it look like someone who needed help," says the Italian Institute of Technology's Sandini, who's leading the project. "Someone … a little sad."

When I met Movellan he seemed flummoxed by the matter of his robot's facial appearance: Should the features be skeletal or soft-tissue, like Einstein's? He was also pondering whether it would be male or female. "All my robots so far have been girls—my daughter has insisted," he explains. "Maybe it's time for a boy." Later, he and his coworkers asked Hanson to help design a face for the Project One robot, which will be named Diego. The "developmental android" will be modeled after a real child, the chubby-cheeked nephew of a researcher in Movellan's lab.

Though Movellan believes that a human infant is born with very little pre-existing knowledge, even he says it comes with needs: to be fed, warmed, napped and relieved of a dirty diaper. Those would have to be programmed into the robot, which quickly gets complicated. "Will this robot need to evacuate?" says John Watson, a University of California at Berkeley professor emeritus of psychology who is a Project One consultant. "Will the thing need sleep cycles? We don't know."

Others outside the project are skeptical that baby robots will reveal much about human

physically as well as cognitively. "To mimic-infant development, robots are going to have to change their morphology in ways that the technology isn't up to," says Ron Chrisley, a cognitive scientist at the University of Sussex in England. He says realistic human features are usually little more than clever distractions: scientists should focus on more basic models that teach us about the nature of intelligence. Human beings learned to fly, Chrisley notes, when we mastered aerodynamics, not when we fashioned realistic-looking birds. A socially capable robot might not resemble a human being anymore than an airplane looks like a sparrow.

Maybe the real magic of big-eyed, round-faced robo-babies is their ability to manipulate our own brains, says Hamid Ekbia, a cognitive science professor at Indiana University and the author of *Artificial Dreams: The Quest for Non-Biological Intelligence*. Infantalized facial features, he says, primarily tap into our attraction to cute kids. "These robots say more about us than they do about machines," says Ekbia. "When people interact with these robots, they get fascinated, but they read beneath the surface. They attribute qualities to the robot that it doesn't have. This is our disposition as human beings: to read more than there is."

Of course, Movellan would counter that such fascination is, in Project One's case, quite essential: to develop like a real child, the machine must be treated like one.

Each Project One researcher defines success differently. Some will declare victory if the robot learns to crawl or to identify basic objects. Watson says he would be grateful to simulate the first three months of development. Certainly, no one expects the robot to progress at the same rate as a child. Project One's timeline extends over four years, and it may take that long before the robot is exposed to people outside the lab—"caregivers" (read: undergrads) who will be paid to baby-sit. Lacking a nursery, the robot will be kept behind glass on a floor beneath Movellan's lab, accessible, for the time being, only to researchers.

As for Movellan, he hopes that the project will "change the way we see human development and bring a more computational bent to it, so we appreciate the problems the infant brain is solving." A more defined understanding of babies' brains might also give rise to new approaches to developmental disorders. "To change the questions that psychologists are asking—that to me is the dream," Movellan adds. "For now it is, how do you get its arm to work, the leg to work? But when we put the pieces together, things will really start to happen."

Before leaving the lab, I stop to bid goodbye to Einstein. All is not well with the robot. Its eye cameras have become obsessed with the glowing red exit sign over the workshop's door. Hanson switches the robot off and on; its movements are palsied; its eyes roll. Its German accent isn't working and the tinny-sounding conversational software seems to be on the fritz. Hanson peers into its eyes. "Hi there," he says. "Can you hear me? Are you listening?"

Einstein: (No response.)
Hanson: Let's get into the topic of compassion.
Einstein: I don't have good peripheral vision.
Einstein: (Continuing.) I am just a child. I have a lot to learn, like what it is to truly love.

Students working nearby are singing along to a radio blasting Tina Turner's "What's Love Got to Do With It," oblivious to Einstein's plight. For me, though, there is something almost uncomfortable about watching the robot malfunction, like seeing a stranger struggle with heavy suitcases. Does this count as magic?

On a worktable nearby, something catches my eye. It is a copy of a Renaissance-era portrait of Mary and the infant Jesus—Carlo Crivelli's *Madonna con Bambino*, the engineers say, which another robot in the room is using to practice analyzing images. The painting is the last thing I expect to see among the piles of tools and snarls of wires, but it occurs to me that building a humanoid robot is also a kind of virgin birth. The child in the painting is tiny but already standing on its own. Mary's eyes are downcast and appear troubled; the baby stretches one foot forward, as though to walk, and gazes up.

2.2.2 ETHICS IN ADVANCED ROBOTICS

FIORELLA OPERTO

In the 1940s, long before we had the technical ability to create autonomous robots, science fiction author Isaac Asimov presciently formulated his Three Laws of Robotics:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Today, we actually have robots that are capable of acting autonomously, that is, that are capable of learning and modifying their own behavior without explicit direction from their human designers. The existence of such machines, and their anticipated introduction into human society, has prompted increased attention to a new field of applied ethics known as roboethics.

Among the kinds of ethical, legal, and social questions that roboethics attempts to address are: How can autonomous robots be programmed to ensure that they behave ethically with respect to recognized human rights? Who is responsible if a robot misbehaves in ways that harm human beings—the designer, the manufacturer, or the user? Is it theoretically possible to regard advanced robots as moral agents in their own rights? And, what are the ethical norms that ought to govern future research and development in the field of robotics?

In this reading, the author, the founder and president of the School of Robotics in Genoa, Italy, explores these questions and discusses some of the answers that roboethicists are currently proposing. While the field is still quite new, one thing is already clear, and that is that Asimov's Three Laws are not going to be sufficient.

FOCUS QUESTIONS

1. What does the author suggest as one reasonable way to address unique situations that might have a number of consequences, some of which are less than desirable?
2. Discuss the author's statement that it is possible a robot could be "more ethical than human beings." Do you agree with that position? Why or why not? What might be evidence of its accuracy?
3. Do you believe the concept of an "autonomous machine" is consistent with the assumption that "a robot will always decide ethically because it cannot behave otherwise," unlike a human? Why or why not?
4. Discuss a few of the "deeper social values and interests" mentioned by the author, which influence "institutional practices in science and technology."

Source: *IEEE Robotics and Automation Magazine* by Institute of Electrical and Electronics Engineers; IEEE Robotics And Automation Society. Reproduced with permission of Institute of Electrical and

KEYWORDS

AI, automatic ethical and regulatory gap, autonomous machine, bionics systems, biorobotics, ELS, HRI, machine learning, operational artificial ethics, prescriptive ethical theory, roboethics, robot rights, service (personal) robots, Universal Declaration of Human Rights

ELS ISSUES IN ADVANCED ROBOTICS BY FIORELLA OPERTO

> The meanings of ethical concepts and rules, in a given situation, should be clear and unambiguous. If they are not, one must undertake to clarify their meanings to the extent possible (. . .) New ethical judgments and cases should be assimilated, where possible, into the existing body of cases, rules, laws, policies, and practices.
>
> Terrell Ward Bynum on Norbert Wiener

Since 2003, the ethical, legal, and societal issues (ELS) in advanced robotics have attracted the increasing and lively interest of academic and professional circles. A similar, although more occasional, debate has also spread to the general public, stimulated either by the novel statements of researchers about recent advancement in robotics or by new and sensitive robotics applications. This increase of contributions and interest occurred hand in hand with the rapid development of research and applications in the service (personal) robot sector, marking the end of the robot segregation era [1].

Many are the current social motivations for a high demand for personal robots. On the economic side, the transitional conclusion of a bull cycle for industrial robotics; on the social side, the new demand for dependable and safety automatic autonomous machines to be employed in human assistance. Although the so-called social robots have not yet been sold as standardized market products for consumers, we can foresee that the research and application will point in this direction, driven both by the market demand and the challenges and richness of technological and scientific issues posed to the roboticist researchers (see the concept of social caretaking [2]).

Some major segments of our lives already depend on sophisticated machines. However, in some critical instances (where robots are entrusted with human lives, such as in medicine, human care, or war theaters), we are approaching an ethical and regulatory gap, primarily because of the lack of new criteria to ascribe responsibility to software agents and learning machines.

In fact, the recent thrust in research and applications programs aiming toward developing robots that are able to cooperate with humans—that are, software agents and robots entrusted with high learning and decision-making functions—presents many theoretical and practical cases in which moral and legal aspects of the responsibility-attribution problem for learning robots may soon become an urgent issue [3]. The responsibility-ascription problem is a central issue in human-robot interaction (HRI).

These and other epistemic limitations and gaps in our current moral codes and regulatory guidelines, concerning software agents, learning machines, and robots, required, in the view of roboticist Gianmarco Veruggio—who coined the word Roboethics and started the debate on this issue—an effort of analysis and original research in this new field of ethics in science and technology.

Roboethics is the applied ethics developed for 1) the identification and analysis of ethical issues that arise from current and upcoming robotics applications and 2) the possible definition of some general guidelines on the issues mentioned earlier. In a broader meaning of the word, roboethics not only applies to the so-called negative rights (that is, the prohibition of some actions) but also the identification of some trends of robotics research and development that enhance what are viewed as human positive rights (the encouragement to some actions): robotics research and applications that promote and enhance human rights for well-being, education, right of medical care, food, housing, and employment. Actually, robotics and bionics systems may potentially make a significant contribution to the

solution of many open human issues. Assistive robots could promote the right of people to live a life of independence and social participation; robotic systems in medicine could foster fundamental human rights by improving the quality of medical operations and protecting the patient's physical integrity; biorobotics have the potential to provide effective therapeutic means to restore lost motor functionalities; and robots for the environment could be extremely useful in the cleaning of polluted areas.

Following the first International Symposium on Roboethics [5], one of the first steps in roboethics was to select a methodology for the identification and analysis of technoethical issues in robotics. To our knowledge, three European projects (which also availed themselves of international experts as well as background studies in information and computer ethics and bioethics) represent the most structured efforts in this field. These are as follows: the Euron Roboethics Atelier, 2006 [6]; Ethicbots 2006–2008 [7]; and Coordination Action for Robotics in Europe (CARE) 2008 [8]. Here, the analyzers could 1) identify the most universally shared and enforced declarations, conventions, and agreements about human rights; 2) achieve a shared knowledge about the ethical notions involved in the potential violation and promotion of human right in relation to human interaction with robotic, bionic, and artificial intelligence (AI) systems; and 3) apply their contextual understanding of these ethical notions to some technoethical cases. This article has adopted the same methodology.

The further step was to identify the human rights to be protected and promoted by roboethics, listed and discussed in the Universal Declaration of Human Rights (1948) and the Treaty of Lisbon (2000). These are as follows:

- respect and protection of human dignity and privacy
- right for physical and integrity of the person
- right for liberty and security
- right to protection of personal data
- right for the elderly to lead a life of dignity and independence and to participate in the social and cultural life
- right to integration of persons with disabilities
- fair access to technological resources and social and cultural discrimination (per ages, gender, and census).

A further phase resulted in applying roboethics analysis to the following cases:

Human dignity and privacy: In what way could robotics, in all its applications, affect human functions, capabilities, and rights related to human data protection and privacy issues? On the side of positive rights, how could these technologies suggest a further definition of the concepts of human liberty, dignity, and identity?

Preservation of human identity and transhumans: In the field of bionics and robotics prosthetics, what is the border between restoration and enhancement? What is the human identity that robotics implants should preserve?

Liability and responsibility issues: Who is responsible for the possible malfunctions and damages by autonomous robots to humans and/or property? (issues in the domain of AI and law. This subject also implies the identification of new possible additions to the definition of personhood and agentive capacities and the responsibilities discharged regarding robots).

Psychological effects: The effects of personal robots on human logical and emotional structures and on human relationships (psychology, sociology, HRI, neurosciences, and law).

Cost-benefit analysis: The comparison between robotics applications and all other possible alternatives, evaluating whether or not to add the ethical element among the variables listed in the cost-benefit analysis. This subject also comprises the issues of human replaced by robot, job displacement, and analysis of new educational requirements and professional qualification of the human workforce and operators (economy, public policy, engineering economics, and corporate social responsibility).

Initially, the fields of robotics that are going to raise more urgent and entangled ethical issues are biorobotics and robotics military applications.

The reason is that these two sectors have, with their research and prototypal applications, immediately and directly intervened in the core of the existing corpus of ethical principles, regulations, and laws related to the most sensitive issues of human life. In less dramatic but not less-important instances, assistive and educational robotics has also given rise to some ethical concerns in the field of HRI.

Because of the rich and complex debate in roboethics and the far-reaching scenarios that could develop over the following decades, this article (with the partnership of the aforementioned European projects) adopted a triaging work methodology, analyzing the issues that had the following elements:

Novelty: Issues that have never been coped with; the absentia legis and the lack of regulations, in many cases (bionics and military robotics), underpin a severe responsibility gap.

Emerging: Issues arising in a nonprogrammed way, as the prototypal robotics products are the results of the drive of different instances: research and business.

Complexity: Issues lying in the crossroads of several disciplines (robotics, AI, moral philosophy, psychology, anthropology, and law).

Social pervasiveness: Issues related to current and yet-to-be-released robotics products.

ETHICS IN COMPLEX TECHNOLOGICAL SOCIETIES

For the purpose of this study, we will provide the various meanings of the term *ethics* within the context of modern debate. Traditionally, ethics is the philosophical or theological subject that studies human behavior and assessment criteria for human behaviors and choices. Modern ethics have developed from various points of view along classical philosophy. In the last decades, in our highly complex societies characterized by technoscientific development, the attribution of moral and practical responsibility is becoming more and more difficult with regard to unintentional or collateral consequences of actions and operations that are produced by unidentifiable decisions (group decisions, complex administrative structure, distribution of responsibilities, and computerized operations). Often, one finds it impossible to attribute final responsibility to a single person or to a defined social entity. In the absence of a definition and a precise analysis of the responsibility chain, technologically advanced societies have shifted the issue onto the concept of risk assessment, thereby attributing value to the damage produced by entities that are seemingly devoid of responsibility. In our case, the question is: Who is responsible for any damage that may be caused by an autonomous robot? Is it the designer, manufacturer, programmer, or final user? Often, it will be difficult to obtain easy answers to this question.

From the viewpoint of ethical theories, we observe that individuals possess the ethics of common sense, which provide them with the moral guidelines for decision making, from the small ones to important life decisions. This affects the day-to-day life as well as our actions in our close-knit social relations. Individuals often adopt different ethical theories from time to time: utilitarianists for some decisions, in other situations, or moments generally go by the terms of their upbringing, customs, tradition, and religion (descriptive ethics). These seem to be enough for them.

However, this line of thought reaches its limits when, as social figures or in our profession, we are faced with complex problems, in which our actions may have multiple consequences that are difficult to follow and predict; and when our common sense is faced with the problems we've never dealt with before—for example, bioethical dilemmas. In these situations, the ethics of common sense leads to various paradoxes: we find ourselves without any conceptual resources and in the difficult position of having to pass judgment. In these cases, we need a logical-critical set of ethics (critical ethics) that 1) reveals the implicit and, perhaps, never uncovered assumptions in our ethics of common sense and in our outdated ethical theoretical base and 2) analyses the reasons, the pros and cons, and their origin. Inevitably—we may not even realize it—we resort to prescriptive ethics and ethical theories.

In practice, when faced with general and complex topics, we may refer to the fundamental, relevant values involved in our dilemmas; we may adhere to the more updated morality applied to issues close to ours; or we try to step up from the universally shared prescriptions toward new ethical frontiers. Prescriptive ethical theory, which develops and justifies the principles of moral actions, refers to ethical theories and related guidelines. These in turn represent the general ideas that contain ethical principles to reach an internal and systematic coherence. In defense of the ethical principle, the ethical philosophers either implicitly or explicitly refer to an ethical theory.

In the 20th century, because of the dissatisfaction raised by the issues that are deemed to be the limits to traditional ethical theory—utilitarianism and deontologism inspired by Kant—ethics became fragmented into several forms: rights ethics, virtue ethics, feminist ethics, and applied ethics [9].

For instance, according to virtue ethics—which have less to do with single actions, but more to do with different life styles and ways of life—the most important moral obligation is our personal relationship with the action of making or using robots and the analysis of the effect of robots with respect to the concept of good and happiness: Which robots could contribute most to the happiness to the full and complete quality of life of human beings? [10]. In another instance, considering the need to account for values that are part of a collective conscience, the so-called rights ethics sees human rights to be the most relevant elements in common for a system of cultural and moral pluralism and as being reasonably deemed to be the final manifestation of universal ethics.

One of the challenges to traditional ethical theories by some thinkers regarded the limitation of moral considerations to the problems of humans and their relationship in human society, regardless of other living organisms and environment. Therefore, ethical theories and the new applied ethical theories, in particular, were called upon to include nonhuman entities in their analysis (animal ethics, environmental ethics, and planetary ethics known as the *Gaia theory*) as well as human products (bioculture, computer ethics, and roboethics).

The various applied ethical theories are, in turn, connected and intertwined with other disciplines, including law, sociology (descriptive ethics), economics, and various scientific fields. One of the central themes of applied ethics is the concept of responsibility, which is a moral notion. Legal responsibility determines the rules of the relevant prescriptive ethics, meaning the group of commands and prohibitions adopted by a society or group and that also defines professional ethics. For the purpose of this article (to analyze the human–robot relationship), we shall consider the two main meanings of the term responsibility:

1. the analysis of the identity of the agent of the cause of certain actions and their effects (utilitarianism or consequentialism or teleologism)
2. an expression of motivations that leads an agent to act in a certain way (deontological ethics or Kantian ethics), according to which the individual assesses the consequences of his or her actions.

In the last century, we know that the questions "What authority and what set of moral rules am I obliged to be accountable to? State law? God?" forced many answers (utilitarian or deontological) to a crisis point. Often, the moral response of the individual as well as the moral evolution of our society have led to an opposition to the responsibility toward the nation, church, or traditional roles of social institutions (see Max Weber's ethics of intentions or ethics of individual conscience).

Some serious events related to World War II changed the notion of responsibility and differentiation of roles (i.e., engineers deal with engineering, doctors with medicine, soldiers obey orders from superiors, etc.) and also some famous legal cases, e.g., the Adolph Eichmann trial, during which he defended himself by stating that he was obeying the orders under his administrative responsibilities and that he did not decide on, nor was he aware of, the entire project of Jewish extermination or the debate that ensued after Hiroshima and Nagasaki 1945.

Furthermore, there are different cases in which our personal moral clashes with those adopted and imposed by the society we live in and whose law we live under. Such is the case, for example, of the death penalty (which is in force in various states

in the United States, although contested by internal currents), animal rights, abortion, or euthanasia. In these cases, the implicit moral philosophy in state laws does not coincide with the feelings of many groups of the relevant societies, which leads to this sort of conflict.

Therefore, we observe that, in contemporary societies, the notion of responsibility is not limited to the moral consideration of actions of the agent or cause and deals with the needed conformity of the action with a group of duties (therefore, the analysis of consequences is less decisive). In today's society, just the elements of complexity and technological ruling determines the aspect known as *heterogenesis of ends*, according to which our actions may have consequences that are extremely difficult to estimate, which may even be opposite to our intentions. According to Morin's ecology of action [11], once the action departs from the individual, it lives a life of its own, and combines itself with the environmental conditions (social models, actions, and reactions of other agents), and the final result is beyond the agent's predictive abilities.

Faced with this vacancy in the attribution of individual responsibility, some researchers attempted to identify collective shared responsibilities, were it to be impossible or vain to identify an individual responsibility. In the case of scientific research, science and technology studies (S&TS) expert René von Schomberg proposes to adopt an assessment system based on foresight and knowledge (foresight and knowledge assessment). The author sustains that, because the definition of responsibility is considerably more arduous to define in scientific fields, due to the unintentional consequences, uncertainty, or ignorance of results, instead of identifying the ethical responsibilities post hoc, it is necessary to establish the ethics of the overlap of knowledge between different areas beforehand (synergy: scientists, politicians, etc.), because the quality of knowledge shall determine the ethical value of the applications that will follow.

At the same time, one must constantly ensure that maximum precision of predictions to identify both the wholesomeness of research and relevant applications as well as the potential ethical problems [12].

Other authors have emphasized the need to avoid the overlap between ethical problems and technical solutions: among the latter, the expert of computer ethics, Abbe Mowshowitz [22] states that

> the seemingly eternal social problems are real enough, but to look for their cause in technique or autonomous technology is both mistaken and harmful. We should not blame technology for human failures. (..) Autonomous technology contributes to the belief in technological determinism, i.e., reinforces belief in the inability of people to make significant choices in their lives. It directs attention away from wielders of power to systems of reified collectivities. The law is smarter than the social sciences—it defines the corporation, for example, as a fictional person for purposes of assigning responsibility and does not absolve key actors of their responsibilities. Institutions should be seen as convenient fictions that help explain individual decision and action (..) Only the actions of human beings can be alienating or dehumanizing. Reification of technology allows for an illegitimate transference of responsibility from persons to a fictional social construct, and at the same time, impedes our ability to come to grips with the very real ethical challenges posed by the uses of technology.

In conclusion, we observe that, in roboethics, the definition of moral responsibility and the resulting notion of liability—which is central in human–robot relationship— could differ according to the philosophical assumptions which, knowingly or unknowingly, have been adopted.

ROBOETHICS OR ROBOT ETHICS?

Roboethics was originally conceived as human-responsibility ethics. The roboticists and ethicists that contributed to their creation highlighted the following aspects:

- problems regarding robot autonomy
- problems related to warfare applications
- problems in human–robot relations (dependence, privacy, robot appearance, and potential confusion between natural and artificial)
- digital divide (for nations, genders, and ages)
- ethical dimension of technology.

According to Veruggio [13],

> roboethics is an applied ethics whose objective is to develop scientific/cultural/technical tools that can be shared by different social groups and beliefs. These tools aim to promote and encourage the development of robotics for the advancement of human society and individuals and to help preventing its misuse against humankind.

This and similar definitions imply that robotics and its applications are subject to moral judgment and human intervention. According to this perspective, roboethics is not artificial ethics and not even the regulatory system of dependability and safety. Roboethics indicated that the individual and society may intervene upstream on the direction of robotics and its products. The individual may, for example, limit the use of robots, according to a wider precautionary principle, in the absence of the necessary safety precautions when missed for humans.

As researchers, individuals may refuse to design robots that are deemed harmful or hazardous; as professional bodies, they may discuss and decide upon an appropriate professional ethic.

As already discussed, roboethics has adopted the principles established by the Charter of Human Rights and the Lisbon Treaty. However, these certainly do not satisfy the realm or depth of the ethical debate. If last century extended the ethical realm to an increasing number of elements, including animals and our planet among those entitled to rights, we witness an increasing application rights ethics to even broader categories, thereby crossing over the human/organic barrier.

In the definition of roboethics, as developed by some authors [6], [14], the agreement with the thesis of American moral and political philosopher John Rawls and to his reflective equilibrium is explicit, according to which the charters and treaties, although advanced in an attempt to associate the highest number of participants and commendable in an attempt to widen the promotion of rights to a greater number of members and functions as possible, do not satisfy the possibility of positive progress, nor can they provide for the entire range of criminal uses of robots.

During the past six years of discussion on roboethics, various positions have emerged regarding the ELS issues in robotics. Just as some authors have continued to use the term roboethics, others have used robot ethics. These two terms don't always indicate different notions. However, robot ethics has been often used to indicate 1) the artificial ethics of robots or the morality of robots and 2) the group of prescriptions, rules, and regulations regarding robot safety and dependability. Although the term robot ethics may be used in the latter sense, because roboethics also studies issues of dependability and safety, the difference between roboethics and robot ethics is more complex.

Some authors see robot ethics as an artificial morality and the robots as autonomous agents. Their theses in this context are different, but they can be linked to this cluster of arguments.

The first states that sophisticated autonomous robots, because they are equipped with intelligence and a certain kind of conscience, they include an ethical system that can learn and evolve or are able to decide between good or evil. According to this position, morality and immorality constitute a gradual continuum: such is the case for children, which, in our society, are not deemed to be fully responsible for their actions or the disabled [15], [16]. As they are quasi-moral agents, robots are also subject to ethical behavior, equipped with some rights of their own.

The second position in robotics morality states that robots, as devoid of emotion and passion, and as they are equipped with rational behavior (because that is how they are programmed), could be more ethical than human beings [17]. These two positions have one aspect in common: they state that, when faced with the complexity of robot technology, it is difficult to assign responsibility for any possible damage they may cause; or when faced with irrational and complex situations (war theaters and other circumstances where an immediate answer—such as a gut reaction—is required), the placing of trust in an automated morality is the only solution. Thus, the focus has shifted from the issue of human ethics and has moved onto operational

artificial ethics, which therefore avoids the debate begun by Sanremo [5] and that was essentially focused on the issue: when is it right to limit the autonomy of robots as final products?

When one tries to shift an ethical problem onto a technical solution, there are some concepts that fall into the background, such as 1) important underlying assumptions about the adopted ethical theories; 2) the relevant glossary, meaning an ambiguous use of terminology. In our case, here the keyword is *machine learning*.

Although searching for new solutions is more than laudable and auspicious—for instances, when faced with the horror of war and above all those of this century—coherence with the philosophy of law imposes that every new form of ethics cannot demean the rights acquired by the previous ones: "Any new ethics must deal with the same substance as the old role responsibility ethics, namely, with values and norms that restrict or delimit human action and thus enable or guide traditional decision making" [12]. This statement means that roboethics and robot ethics should be coherent with the shared sets of rules.

Overall, every thesis in favor of robotics morality per se is based on the assumption of predictability (or rationality and foreseeability) of the behavior of an autonomous robot. A robotic morality provides, in the conception of its authors, that the robot shall behave ethically according to what it has been programmed with and what it has learned. In a word, it cannot deviate from its incorporated laws: unlike a human being that, although supposedly brought up to be good, may decide to do evil unexpectedly, a robot will always decide ethically because it cannot behave otherwise.

Here, the critical analysis should focus on two critical elements: 1) the kind of ethics that the robot expresses and 2) the unpredictability issue in learning machines. The first question implies, in our view, that no robot ethics can avoid the needed, deep and broad debate on the human ethical principles grounding roboethics and their enforcement. The second critical element is that, according various experts [18], [19], "programmers, manufacturers, and users may not be in the position to predict what a learning robot will do in normal operating environments and to select an appropriate course of action on the basis of this prediction" [3]. This means that not only it would be advisable that learning robots should be self-evidently distinguishable from nonlearning robots, but that in the former case, the learning process be made transparent for the robot user.

An ethically accurate analysis of the bonds and limits of the potential sphere of action of a robotics morality should indeed take into careful account the epistemic and logic-hidden faults related to both ethical theories and technical constraints. The fact that the current development of non-segregated robotic systems does not yet allow for a precise modeling of every possible environmental factor or for a final definition of normalcy in rich operating conditions should be the basis for a more prudent discourse on intelligent robots behaving ethically.

Moreover, a careful assessment of the responsibility issues in attributing to robot rights, which belong only to humans, and the analysis of the problems encountered in assessing the environmental factors affecting robotics behavior, would advise to adopt the enlarged version of the precautionary principle, limiting robotics autonomy when all the environmental factors are not precisely assessed.

CONCLUSIONS: AN OPEN DEBATE

From the manifold and depth of the considerations around roboethics over the last six years and considering the new ethical issues involved that have never been broached, it appears as if, beyond the different point of views, one major side question is posed, which would need a more general answer: Which direction robotics is going to take? And which should it take? The hope of the many, whose consideration of the importance of robots in society, goes hand in hand with ethical related concerns, is that a promising alliance between robotics and the field of science and technology studies should also be swiftly established.

It has been widely recognized—although not very often practically accepted—that institutional

practices in science and technology are tacitly shaped and framed by deeper social values and interests [20]. These include

- important political–economic relationship to science and technology
- shifting from science as independent republic to science as cooperating in innovation and in the knowledge economy
- impacts of the increasing commercialization of science in particular areas affecting public trust
- loss of credibility and senses of unease in the general public [21].

The researches in science and society developed by social scientist Sheila Jasanoff and collaborators have sustained that technoscientific knowledge stabilizes in society through a complex and articulated process of negotiation and then very seldom experts' opinions are exempt from uncovered assumptions.

It should be clear to all the parties involved in the process of molding roboethics that this free, open, and relatively untroubled debate is possible, because until now no dramatic incident in the field occurred. And, we hope it never will. Robotics research is driven by future scientific visions, and it so should be. Limitations of the freedom of scientific research should be very carefully discussed and almost never imposed. Limitations should be painstakingly decided on for marketed robotics applications.

However, some from the robotics industry have already expressed their concern that, at the first sense of unease from some social groups toward new robotics applications, some limits will be decided on. Increasingly, if a dramatic incident was to occur, the extraordinary character of the situation will impose severe constraints. We have already witnessed those occurrences in bioethics (stem cells, etc.). Among the recommendations resulting from the analysis of ELS issues in robotics developed in the frame of CARE, there is a need to "avoid that ELS issues in robotics could become a barrier to further progress of our field."

If some incidents were to occur, the outdoor process of discussion will be channeled in the indoor environment of the experts committee, relegating the other parties (roboticists, ethicists, stakeholders, and society as a whole) to the role of concerned observers.

For these reasons, we feel that all people concerned with roboethics should take vantage from this fortunate window of fresh and free debate, defining with careful consideration and wise temperance of language the general ethical assessments and rules for future robotics.

REFERENCES

1. B. Siciliano and O. Khatib, Springer Handbook of Robotics. Berlin Heidelberg: Springer-Verlag, 2008.
2. C. Breazeal, "Emotion and sociable humanoid robots," Int. J. Human Comput. Stud, vol. 59, no. 1–2, pp. 119–115, July 2003.
3. G. Tamburrini and D. Marino, "Learning robots and human responsibility," Int. Rev. Inform. Ethics, vol. 6, pp. 46–52, Dec. 2006.
4. G. Veruggio, "Io robotico," Le Scienze (Italian Edition of Scientific American), no. 434, pp. 40–47, Oct. 2004.
5. G. Veruggio, (2005, Apr. 18). The birth of roboethics. presented at Proc. IEEE Int. Conf. Robotics and Automation Workshop on Roboethics (ICRA 2005), Barcelona [Online]. Available: http://www.roboethics.org/icra2005/veruggio.pdf
6. G. Veruggio, (2006, Dec. 4–6). The Euron Roboethics Atelier. For the Euron Roboethics Roadmap presented at Proc. 6th IEEE-RAS Int. Conf. Humanoid Robots, Genoa, Italy [Online]. Available: http://www.euron.org/activities/projects/roboethics
7. ETHICBOTS project. Emerging Technoethics of Human Interaction with Communication, Bionic, and robOTic systems. (2005, Nov. 1–2008, Apr. 30). Coordination action, FP6—Science and society [Online]. Available: http://ethicbots.na.infn.it/
8. CARE. Coordination Action for Robotics in Europe. The strategic research agenda. European Robotics Technology Platform, [Online]. Available: http://www.robotics-platform.eu/cms/index.php?idcat=26
9. A. Fabris, Quale etica per la robotica? in Ethicbots, Etica e Robotica, (Teoria, Rivista di filosofia). Pisa: Edizioni ETS, 2007, pp. 7–17.
10. M. Nussbaum and A. Sen, The Quality of Life. Oxford: Clarendon, 1993.

11. E. Morin, La Methode. Ethique: Editions de Seuil, 2004.
12. R. von Schomberg, "From the ethics of technology towards an ethics of knowledge policy: Implications for robotics," AI Soc., vol. 22, no. 3, pp. 331–345, 2007.
13. G. Veruggio and F. Operto, "Ethical, legal and societal issues in the strategic research agenda of the coordination action for robotics in Europe," in Proc. Workshop on Roboethics (ICRA09), Kobe, Japan, May 17, 2009.
14. G. Tamburrini, "Robot ethics: A view from the philosophy of science," in Ethics and Robotics, R. Capurro and M. Nagenborg, Eds. Heidelberg: IOS Press, 2009, pp. 11–23.
15. P. M. Asaro, "What should we want from a robot ethic?," in Ethics and Robotics, M. Capurro and M. Nagenborg, Eds. Heidelberg: IOS Press, 2009, pp. 1–11.
16. W. Wallach and C. Allen, Moral Machines: Teaching Robots Right from Wrong. Oxford, New York: Oxford Univ. Press, 2009.
17. R. Arkin, Governing Lethal Behavior in Autonomous Robots, Boca Raton: CRC Press, 2009.
18. J. Altmann, "Preventive arms control for uninhabited military vehicles," in Ethics and Robotics, R. Capurro and M. Nagenborg, Eds. Heidelberg: IOS Press, 2009, pp. 69–83.
19. C. Tamburrini, "AI and Popper's solution to the problem of induction," in Karl Popper: A Centennial Assessment, vol. 2, Metaphysics and Epistemology, I. Jarvie, K. Milford, and D. Miller, Eds. London: Ashgate, pp. 265–282.
20. M. Callon, M. E. Gonçalves, S. Jasanoff, M. Jepsen, P. B. Joly, Z. Konopasek, S. May, C. Neubauer, A. Rip, K. Siune, A. Stirling, and M. Tallacchini, "Taking European knowledge society seriously," Report of the Expert Group on Science and Governance to the Science, European Commission, Directorate-General for Research, Science, Economy and Society Editor, Brussels, Report no. EUR 22700, 2007.
21. M. Polanyi, Personal Knowledge–Towards a Postcritical Philosophy, London: Routledge, 1958.
22. A. Mowshowitz, "Technology as excuse for questionable ethics," AI Soc., vol. 22, no. 3, pp. 271–282, 2008.
23. D. Bloor, Knowledge and Social Imagery, London: Routledge, 1976.
24. G. Veruggio and K. Abney, "Roboethics: The applied ethics for a new science," in Robot Ethics: The Ethical and Social Implications of Robotics, P. Lin, K. Abney, and G Bekey, Eds. Cambridge, MA: MIT Press, to be published.
25. G. Veruggio and F. Operto, "Roboethics: Social and ethical implications of robotics," in Handbook of Robotics. Heidelberg: Springer-Verlag, 2008, ch. 64.
26. J. H. Moor, "What is computer ethics?," in Computers and Ethics, T. W. Bynum, Ed. Oxford, U.K.: Blackwell, 1985, pp. 266–275, Also in Meta-philosophy, vol. 16, no. 4, Oct. 1985.
27. J. Weber, "Robotic warfare, human rights and the rhetoric of ethical machines," in Ethics and Robotics, R. Capurro and M. Nagenborg, Eds. Heidelberg: IOS Press, 2009, pp. 83–105.
28. L. Floridi, "Information ethics: Its nature and scope," in Information Technology and Moral Philosophy, J. van den Hoven and J. Weckert, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2008, pp. 40–65.
29. M. Nagenborg, R. Capurro, J. Weber, and Ch. Pingel, "Ethical regulations on robotics in Europe," AI Soc, vol. 22, no. 3, pp. 349–366, 2007.
30. N. Sharkey, "The ethical frontiers of robotics," Science, vol. 322, no. 5909, pp. 1800–1801, 2008.
31. S. Gill and G Tamburrini, "Guest editorial," AI Soc., vol. 3, no. 22, pp. 265–270, 2007.
32. S. Jasanoff, The Fifth Branch: Science Advisers as Policymakers. Cambridge, MA: Harvard Univ. Press, 1990.
33. M. Santono, D. Marino, and G Tamburrini, "Learning robots interacting with humans: From epistemic risk to responsibility," AI Soc., vol. 3, no. 22, pp. 301–314, 2007.
34. T. Bynum, W. Maner, and J. Fodor, Eds. Teaching Computer Ethics. New Haven: Research Center on Computing & Society, 1993.
35. T. W. Bynum, "Norbert Wiener and the rise of information ethics," in Information Technology and Moral Philosophy, J. van den Hoven and J. Weckert, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2008, pp. 8–25.

2.2.3 FINAL JEOPARDY

STEPHEN BAKER

The idea of a machine that could think like a human has fascinated us for generations, and has inspired many science fiction stories. With the exponential increase in the power of computers and the growth of the Internet, many now believe that the basic goals of artificial intelligence may be achieved in the near future. While IBM's Deep Blue computer program defeated Grand Master Gary Kasparov in 1997, the system was totally useless for anything but chess. IBM's *Watson* program was built using insights from the Deep Blue development team and soon demonstrated its unique abilities by competing against humans in a special *Jeopardy* championship.

The idea of testing a powerful, custom-made computer against the knowledge capacities of human mind initially seemed farfetched to engineers and scientists at IBM Research. This selection describes the process by which the IBM team approached the seemingly impossible task of producing a machine that could "understand" the complexities of natural language, in the form of *Jeopardy* trivia questions. In addition to being equipped with the computing power of 2,208 processors and 200 million pages of information, Watson had to be taught logic and strategy that would enable it to interpret an array of complicated trivia questions and select the correct answer from among the answers. The culmination of the IBM team's efforts came in 2011, when Watson defeated the two top *Jeopardy* champions in a specially designed tournament.

Although the IBM scientists who worked on the Watson project doubtlessly found the work intellectually stimulating, it is also likely that IBM saw the possibility of applying Watson's technological advances in commercial applications. Watson is likely to be a harbinger of a new era in which machine or artificial intelligence will become increasingly prevalent in daily life.

FOCUS QUESTIONS

1. What are some possible applications of the computational abilities demonstrated by Watson, which might prove to be marketable in the future? How might Watson's unique skills be an advantage or disadvantage?
2. Although the term *artificial intelligence* is used in this selection, explain whether and why you think it is appropriate or inappropriate to refer to Watson as *intelligent*.
3. In what ways was the process used by Watson in answering *Jeopardy* questions similar to the functioning of a human mind, and how was it different?
4. Do you believe that Watson portends future computers or other machines that will eventually compete with and challenge humans? Why or why not?

Source: "Introduction" from *Final Jeopardy: Man Vs. Machine and the Quest to Know Everything* by Stephen Baker. Copyright © 2011 by Stephen Baker. Reprinted by permission of Houghton Mifflin Harcourt Publishing Company. All rights reserved.

KEYWORDS

algorithms, artificial intelligence, avatar, cognition, cognitive apparatus, computational linguistics, correlate, Daily Double, Deep Blue, factoids, game strategy, key words, latency, level of confidence, machine logic, processors, sentient being, supercomputer

Watson paused. The closest thing it had to a face, a glowing orb on a flat-panel screen, turned, from forest green to a dark shade of blue. Filaments of yellow and red streamed steadily across it, like the paths of jets circumnavigating the globe. This pattern represented a state of quiet anticipation as the supercomputer awaited the next clue. It was a September morning in 2010 at IBM Research, in the hills north of New York City, and the computer, known as Watson, was annihilating two humans, both champion players, in practice rounds of *Jeopardy*! Within months, it would be playing the game on national television in a million-dollar man vs. machine match against two of *Jeopardy*'s all-time greats.

As Todd Crain, an actor and the host of these test games, started to read the next clue, the filaments on Watson's display began to jag and tremble. Watson was thinking—or coming as close to it as a computer could. The $1,600 clue, in the category The Eyes Have It, read; "This facial ware made Israel's Moshe Dayan instantly recognizable worldwide."

The three players—two human and one electronic—could read the words as soon as they appeared on the big *Jeopardy* board. But they had to wait for Crain to read the entire clue before buzzing. That was the rule. As the host pronounced the last word, a light would signal that contestants could buzz. The first to hit the button could win $1,600 with the right answer—or lose the same amount with a wrong one. (In these test matches, they played with funny money.)

This pause for reading gave Watson three or four seconds to hunt down the answer. The first step was to figure out what the clue meant. One of its programs promptly picked apart the grammar of the sentence, identifying the verbs, objects, and key words. In another section, research focused on Moshe Dayan. Was this a person? A place in Israel? Perhaps a holy site? Names like John and Maria would signal a person. But Moshe was more puzzling.

During these seconds, Watson's cognitive apparatus—2,208 computer processors working in concert—mounted a massive research operation through thousands of documents around Moshe Dayan and his signature facial ware. After a second or so, different programs, or algorithms, began to suggest hundreds of possible answers. To us, many of them would look like wild guesses. Some were phrases that Dayan had uttered, others were references to his military campaigns and facts about Israel. Still others cited various articles of his clothing. At this point, the computer launched its second stage of analysis, figuring out which response, if any, merited its confidence. It proceeded to check and recheck facts, making sure that Moshe Dayan was indeed a person, an Israeli, and that the answer referred to something he wore on his face.

A person looking at Watson's frantic and repetitive labors might conclude that the player was unsure of itself, laughably short on common sense, and scandalously wasteful of computing resources. This was all true. Watson barked up every tree from every conceivable angle. The pattern on its screen during this process, circles, exploding into little stars, provided only a hint of the industrial-scale computing at work. In a room behind the podium, visible through a horizontal window, Watson's computers churned, and the fans cooling them roared. This time, its three seconds of exertion paid off. Watson came up with a response, sending a signal to a mechanical device on the podium. It was the size of a large aspirin bottle with a clear plastic covering. Inside was a *Jeopardy* buzzer. About one one-hundredth of a second later, a metal finger inside this contraption shot downward, pressing the button.

Justin Bernbach, a thirty-eight-year-old airline lobbyist from Brooklyn, stood to Watson's left. He had pocketed $155,000 while winning seven straight *Jeopardy* matches in 2009. Unlike Watson, Bernbach understood the sentence. He knew precisely who Moshe Dayan was as soon as he saw the clue, and he carried an image of the Israeli leader in his mind. He gripped the buzzer in his fist and frantically pressed it four or five times as the light came on.

But Watson had arrived first.

"Watson?" said Crain.

The computer's amiable male voice arranged the answer, as *Jeopardy* demands, in the form of a question: "What is eye patch?"

"Very good," Crain said. "An eye patch on his left eye. Choose again, Watson."

Bernbach slumped at his podium. This match with the machine wasn't going well.

It was going magnificently for David Ferrucci. As the chief scientist of the team developing the *Jeopardy* computer, Ferrucci was feeling vindicated. Only three years earlier, the suggestion that a computer might match wits and word skills with human champions in *Jeopardy* sparked opposition bordering on ridicule in the halls of IBM Research. And the final goal of the venture, a nationally televised match against two *Jeopardy* legends, Ken Jennings and Brad Rutter, seemed risky to some, a bit déclassé to others. *Jeopardy*, a television show, appeared to lack the timeless cachet of chess, which IBM computers had mastered a decade earlier.

Nonetheless, Ferrucci and his team went ahead and built their machine. Months earlier, it had fared well in a set of test matches. But the games revealed flaws in the machine's logic and game strategy. It was a good player, but to beat Jennings and Rutter, who would be jousting for a million-dollar top prize, it would have to be great. So they had worked long hours over the summer to revamp Watson. This September event was the coming-out party for Watson 2.0. It was the first of fifty-six test matches against a higher level of competitor: people, like Justin Bernbach, who had won enough matches to compete in

In these early matches, Watson was having its way with them. Ferrucci, monitoring the matches from a crowded observation booth, was all smiles. Keen to promote its *Jeopardy* phenom, IBM's advertising agency, Ogilvy & Mather, had hired a film crew to follow Ferrucci's team and capture the drama of this opening round of championship matches. The observation room was packed with cameras. Microphones on long booms recorded the back-and-forth of engineers as they discussed algorithms and Watson's response time, known as latency. Ferrucci, wearing a mike on his lapel, gave a blow-by-blow commentary as Watson, on the other side of the glass, strutted its new and smarter self.

It was almost as if Watson, like a person giddy with hubris, was primed for a fall. The computer certainly had its weaknesses. Even when functioning smoothly, it would make its share of wacky mistakes. Right before the lunch break, one clue asked about "the inspiration for this title object in a novel and a 1957 movie [which] actually spanned the Mae Khlung." Now, it would be reasonable for a computer to miss "The Bridge over the River Kwai," especially since the actual river has a different name. Perhaps Watson had trouble understanding the sentence, which was convoluted at best. But how did the computer land on its outlandish response, "What is Kafka?" Ferrucci didn't know. Those things happened, and Watson still won the two morning matches.

It was after lunch that things deteriorated. Bernbach, so frustrated in the morning, started to beat Watson to the buzzer. Meanwhile, the computer was making risky bets and flubbing entire categories of clues. Defeat, which had seemed so remote in the morning, was now just one lost bet away. It came in the fourth match. Watson was winning by $4,000 when it stumbled on this Final Jeopardy clue: "On Feb. 8, 2010, the headline in a major newspaper in this city read: 'Amen! After 43 years, our prayers are answered.'" Watson missed the reference to the previous day's Super Bowl, won by the New Orleans Saints. It bet $23,000 on Chicago. Bernbach also botched the clue, guessing New York. But he bet less than Watson, which made him the first person to

In the sixth and last match of the day, Watson trailed Bernbach, $16,200 to $21,000. The computer landed on a Daily Double in the category Colleges and Universities, which meant it could bet everything it had on nailing the clue. A $5,000 bet would have brought it into a tie with Bernbach. A larger bet, while risky, could have catapulted the computer toward victory. "I'll take five," Watson said.

Five. Not $5,000, not $500. Five measly dollars of funny money. The engineers in the observation booth were stunned. But they kept quieter than usual; the cameras were rolling.

Then Watson crashed. It occurred at some point between placing that lowly bet and attempting to answer a clue about the first Catholic college in Washington, D.C. Watson' "front end," its voice and avatar, was waiting for its thousands of processors, or "back end," to deliver an answer. It received nothing. Anticipating such a situation, the engineers had prepared set phrases. "Sorry," Watson said, reciting one of them, "I'm stumped." Its avatar displayed a dark blue circle with a single filament orbiting mournfully in the Antarctic latitudes.

What to do? Everyone had ideas. Maybe they should finish the game with an older version of Watson. Or perhaps they could hook it up to another up-to-date version of the program at the company's Hawthorne labs, six miles down the road. But some worried that a remote connection would slow Watson's response time, causing it to lose more often on the buzz. In the end, as often happens with computers, a reboot brought the hulking *Jeopardy* machine back to life. But Ferrucci and his team got an all-too-vivid reminder that their *Jeopardy* player, even as it prepared for a debut on national television, could go haywire or shut down at any moment. When Watson was lifted to the podium, facing banks of cameras and lights, it was anybody's guess how it would perform. Watson, it was clear, had a frighteningly broad repertoire.

Only four years earlier, in 2006, Watson was a prohibitive long shot, not just to win at *Jeopardy* but even to be built. For more than a year, the head of IBM Research, a physicist named Paul Horn, had been pressing a number of teams at the company to pursue a *Jeopardy* machine. The way he saw it, IBM had triumphed in 1997 with its chess challenge. The company's machine, Deep Blue, had defeated the reigning world champion, Garry Kasparov. This burnished IBM's reputation among the global computing elite while demonstrating to the world that computers could rival human beings in certain domains associated with intelligence.

That triumph left IBM's top executives hungry for an encore. Horn felt the pressure. But what could the researchers get a computer to do? Deep Blue had rifled through millions of scenarios per second, calculated probabilities, and made winning moves. It had aced applied math. But it had skipped the far more complex domain of words. This, Horn thought, was where the next challenge would be. Far beyond the sixty-four squares on a chess board, the next computer should charge into the vast expanse of human language and knowledge. For the test, Horn settled on *Jeopardy*, which debuted in 1964 and now attracted some nine million viewers every weeknight. It was the closest thing in the United States to a knowledge franchise. "People associated it with intelligence," Horn later said.

There was one small problem. For months, he couldn't get any takers. *Jeopardy*, with its puns and strangely phrased clues, seemed too hard for a computer. IBM was already building machines to answer questions, and their performance, in speed and precision, came nowhere close to that of even a moderately informed person. How could the next machine grow so much smarter?

And while researchers regarded the challenge as daunting, many people, Horn knew, saw it precisely the other way. Answering questions? Didn't Google already do that?

Horn eventually enticed David Ferrucci and his team to pursue his vision. Ferrucci, then in his midforties, wore a dark brown beard wrapped around his mouth and wire-rimmed glasses. An expert in Artificial Intelligence (AI), he had a native New Yorker's gift of the gab and an openness, even about his own life, that was at times jolting. ("I have a growing list of potentially mortal diseases," he said years later. "People order an MRI a week for me.") But he also had a wide and

ranging intellect. Early in his tenure at IBM he and a friend tried, in their spare time, to teach a machine to write fiction by itself. They trained it for various literary themes, from love to betrayal, and they named it Brutus, for Julius Caesar's traitorous comrade. Ferrucci was comfortable talking about everything from the details of computational linguistics to the evolution of life on earth and the nature of human thought. This made him an ideal ambassador for a *Jeopardy* machine. After all, his project would raise a broad range of issues, and fears, about the role of brainy machines in society. "Would they compete for jobs? Could they establish their own agendas, like the infamous computer HAL, in *2001: A Space Odyssey*, and take control? What was the future of knowledge and intelligence, and how would brains and machines divvy up the cognitive work? Ferrucci was always ready with an opinion. At the same time, he could address the strategic questions—how these machines would fit into hundreds of businesses, and why the project he was working on, as he saw it, went far beyond Google.

The Google question was his starting point; until people understood that his machine was not just a souped-up search engine, the project made little sense. For certain types of questions, Ferrucci said, a search engine could come up with answers. These were simple sentences with concrete results, what he and his team called factoids. For example: "What is the tallest mountain in Africa?" A search engine would pick out the three key words from that sentence and in a fraction of a second suggest Kenya's 19,340-foot-high Kilimanjaro. This worked, Ferrucci said, for about 30 percent of *Jeopardy* questions. But performance at that low level would condemn Watson to defeat at the hands of human amateurs.

A *Jeopardy* machine would have to master far thornier questions. Just as important, it would have to judge its level of confidence in an answer. Google's algorithms delivered users to the statistically most likely outposts of the Web and left it to the readers to find the answers. "A search engine doesn't know that it understood the question and that the content is right," Ferrucci said. But a *Jeopardy* machine would have to find answers and then decide for itself if they were worth betting on. Without this judgment, the machine would never know when to buzz. It would require complex analysis to develop this "confidence."

Was it worth it? Didn't it make sense for machines to hunt through mountains of data and for people to rely on their exquisitely engineered brains to handle the final judgments? This seemed like a reasonable division of labor. After all, processing language and spotting answers come easily to humans and are so hard for machines.

But what if machines could take the next step? What if they could go beyond locating bits and pieces of information and help us to understand it? "I think there are 1.4 million books on sale online," Ferrucci said one afternoon. He was sharing a bottle of his own wine, a Shiraz blend that he'd cooked up in the winery attached to his kitchen in the northern suburbs of New York. He was in an expansive mood, which led him to carry out energetic dialogues with himself, asking questions and answering them emphatically. "You can probably fit all the books that are on sale on about two terabytes that you can buy at OfficeMax for a couple hundred dollars. You get every book. Every. Single. Book. Now what do you do? You can't read them all! What I want the computer to do," he went on, "is to read them for me and tell me what they're about, and answer my questions about them. I want this for *all* information. I want machines to read, understand, summarize, describe the themes, and do the analysis so that I can take advantage of all the knowledge that's out there. "We humans need help. I know I do!"

Before building a *Jeopardy* machine, Ferrucci and his team had to carry this vision one step further: They had to make a case that a market existed outside the rarefied world of *Jeopardy* for advanced question-answering technology. IBM's biggest division, after all, was Global Services, which included one of the world's largest consultancies. It sold technical arid strategic advice to corporations all over the world. Could the consultants bundle this technology into their offerings? Would this type of machine soon be popping up in offices and answering customers' questions on the phone?

Ferrucci envisioned a *Jeopardy* machine spawning a host of specialized know-it-alls. With the right training, a technology that could understand everyday language and retrieve answers in a matter of seconds could fit just about anywhere. Its first job would likely be in call centers. It could answer tax questions, provide details about bus schedules, ask about the symptoms of a laptop on the fritz and walk a customer through a software update. That stuff was obvious. But there were plenty of other jobs. Consider publicly traded companies, Ferrucci said. They had to comply with a dizzying assortment of rules and regulations, everything from leaks of inside information in e-mails to the timely disclosure of earnings surprises or product failures to regulators and investors. A machine with Watson's skills could stay on top of these compliance matters, pointing to possible infractions and answering questions posed in ordinary English. A law firm could call on such a machine to track down the legal precedent for every imaginable crime, complaint, or trademark.

Perhaps the most intriguing opportunity was in medicine. While IBM was creating the *Jeopardy* machine, one of the top medical shows on television featured a nasty genius named Gregory House. In the beginning of most episodes a character would collapse, tumbling to the ground during a dance performance, a lovers' spat, or a kindergarten class. Each one suffered from a different set of symptoms, many of them gruesome. In the course of the following hour, amid the medical team's social and sexual dramas, House and his colleagues would review the patient's worsening condition. There had to be a pattern. Who could find it and match it to a disease, ideally before the patient died? Drawing from their own experience, the doctors each mastered a diverse set of data. The challenge was to correlate that information to the ever-changing list of symptoms on the white board in House's office. Toward the end of the show, House would often notice some detail—perhaps a lyric in a song or an unlikely bruise. And that would lead his magnificent mind straight to a case he remembered or a rites in New Guinea. By the end of the show, the patient was headed toward recovery.

An advanced question-answering machine could serve as a bionic Dr. House. Unlike humans, it could stay on top of the tens of thousands of medical research papers published every year. And, just as in *Jeopardy*, it could come up with lists of potential answers, or diagnoses, for each patient's ills. It could also direct doctors toward the evidence it had considered and provide its reasoning. The machine, lacking common sense, would be far from perfect. Just as the *Jeopardy* computer was certain to botch a fair number of clues, the diagnoses coming from a digital Dr. House would sometimes be silly. So people would still run the show, but they'd be assisted by a powerful analytical tool.

In those early days, only a handful of researchers took part in *fat Jeopardy* project at IBM. They could fit easily into Ferrucci's office at the research center in Hawthorne, New York, about thirty-five miles north of New York City (and a fifteen-minute drive from corporate headquarters, in Armonk). But to build a knowledge machine, Ferrucci knew, would require extensive research and development. In a sense, a *Jeopardy* machine would represent an entire section of the human brain. To build it, he would need specialists in many aspects of cognition. Some would be experts in language, others in the retrieval of information. Some would attempt to program the machine with judgment, writing algorithms to steer it toward answers. Others would guide it in so-called machine learning, so that it could train itself to pick the most statistically promising combinations of words and pay more attention to trustworthy sources. Experts in hardware, meanwhile, would have to build a massive computer, or a network of them, to process all of this work. Assembling these efforts on a three-year timetable amounted to a daunting management challenge. The cost of failure would be humiliation, for both the researchers and their company.

Other complications came from the West Coast, specifically the Robert Young building on the Sony lot in Culver City, a neighborhood just south of Hollywood. Unlike chess, a treasure we

Pictures Entertainment, an arm of the Japanese consumer electronics giant. The *Jeopardy* executives, led by a canny negotiator named Harry Friedman, weren't about to let IBM use their golden franchise and their millions of viewers on its own terms. Over the years, the two companies jousted over the terms of the game, the placement of logos, access to stars such as Ken Jennings and Brad Rutter, and the writing of *Jeopardy* clues. They even haggled over the computer's speed on the buzzer and whether IBM should take measures to slow it to a human level. These disagreements echoed until the eve of the match. At one point, only months before the showdown, *Jeopardy*'s executives appeared to be on the verge of pulling the plug on the entire venture. That would have left IBM's answering computer, the product of three intense years of research, scrounging for another game to play. This particular disagreement was resolved. But the often conflicting dictates of promotion, branding, science, and entertainment forged a fragile and uneasy alliance.

The *Jeopardy* project also faced harsh critics within IBM's own scientific community. This was to be expected in a field—Artificial Intelligence—where the different beliefs about knowledge, intelligence, and the primacy of the human brain bordered on the theological. How could there be any consensus in a discipline so vast? While researchers in one lab laboriously taught machines the various meanings of the verb "to do," futurists just down the hall insisted that computers would out-race human intelligence in a couple of decades, controlling the species. Beyond its myriad approaches and outlooks, the field could be divided into two camps, idealists and pragmatists. The idealists debated the nature of intelligence and aspired to build computers that could think conceptually, like human beings, perhaps surpassing us. The pragmatists created machines to carry out tasks. Ferrucci, who had promised to have a television-ready computer by early 2011, fell firmly into the second camp—and his team attracted plenty of barbs for it. The *Jeopardy* machine would sidestep the complex architecture of the brain and contrive to answer questions without truly understanding them. "It's just another gimmick," said Salil Rao, a professor in computer science at MIT who's attempting to teach computers to conceptualize forty-eight different verbs. "It's not addressing any fundamental problems." But as Ferrucci countered, teaching a machine to answer complex questions on a broad range of subjects would represent a notable advance, whatever the method.

IBM's computer would indeed come to answer a dizzying variety of questions—and would raise one of its own. With machines like this in our future, what do we need to store in our own heads? This question, of course, has been recurring since the dawn of the Internet, the arrival of the calculator, and even earlier. With each advance, people have made internal adjustments and assigned ever larger quantities of memory, math, geography, and more to manmade tools. It makes sense. Why not use the resources at hand? In the coming age, it seems, forgoing an effective answering tool will be like volunteering for a lobotomy.

In a sense, many of us living through this information revolution share something with the medieval monks who were ambushed by the last one. They spent years of their lives memorizing sacred texts that would soon be spilling off newfangled printing presses. They could have saved lots of time, and presumably freed up loads of capacity, by archiving those texts on shelves. (No need to discuss here whether the monks were eager for "free time," a concept dangerously close to Sloth, the fourth of the Seven Deadly Sins.) In the same way, much of the knowledge we have stuffed into our heads over the years has been rendered superfluous by new machinery.

So what does this say about Ken Jennings and Brad Rutter, the humans preparing to wage cognitive war with Watson? Are they relics? Sure, they might win this round. But the long-term prognosis is grim. Garry Kasparov, the chess master who fell to IBM's Deep Blue, recently wrote that the golden age of man-machine battles in chess lasted from 1994 to 2004. Before that decade, machines were too dumb; after it, the roles were reversed. While knowledge tools, including Watson, relentlessly advance, our flesh-and-blood brains, some argue, have stayed more or less the same for forty thousand years, meaning

evolutionary water from the Cro-Magnon cave painters to Quentin Tarantino,

A few decades ago, know-it-alls like Ken Jennings seemed to be the model of human intelligence. They aced exams. They had dozens of facts at their fingertips. In one quiz show that predated *Jeopardy*, *College Bowl*, teams of the brainiest students would battle one another for the honor of their universities. Later in life, people turned to them in boardrooms, university halls, and cocktail parties for answers. Public education has been designed, in large part, to equip millions with a ready supply of factual answers. But if Watson can top them, what is this kind of intelligence worth?

Physical strength has suffered a similar downgrade. Not so long ago, a man with superhuman strength played a valuable role in society. He was a formidable soldier. When villagers needed boulders moved or metal bent, he got the call. After the invention of steam engines and hydraulic pumps, however, archetypal strongmen were shunted to jobs outside the productive economy. They turned to bending metal in circuses or playing noseguard on Sunday afternoons. For many of us, physical strength, once so vital, has become little more than a fashion statement. Modern males now display muscles as mating attire, much the way peacocks fan their otherwise useless feathers.

It would be all too easy to dismiss human foes of the IBM machine as cognitive versions of circus strongmen: trivia wunderkinds. But from the very beginning, Ferrucci saw that the game required far more than the simple regurgitation of facts. It involved strategy, decision making, pattern recognition, and a knack for nuance in the language of the clues. As the computer grew from a whimsical idea into a *Jeopardy* behemoth, it underwent an entire education, triumphing in some areas, floundering in others. Its struggles, whether in untangling language or grappling with abstract ideas, highlighted the areas in which humans maintain an edge. It is in the story of Watson's development that we catch a glimpse of the future of human as well as machine intelligence.

The secret is wrapped up in the nature of knowledge is an entire universe, a welter of sensations and memories, desires, facts, skills, songs and images, words, hopes, fears, and regrets, not to mention love. But for those hoping to build intelligent machines, it has to be simpler. Broadly speaking, it falls into three categories: sensory input, ideas, and symbols. Consider the color blue. Sensory perception is the raw material of knowledge. It's something that computers and people alike can perceive, each in their own fashion. Now think of the word "sky." Those three letters are a symbol for the biggest piece of blue in our world. Computers can handle such symbols. They can find references to "sky" in documents and, when programmed, correlate it with others, such as "blue," "clouds," and "heaven." A computer can master both sensory data and symbols. It can count, categorize, search, and store them. But how about this snippet from Lord Byron: "Friendship is love without his wings." That sentence represents the third realm of knowledge: ideas. How can a machine make sense of them? In these early years of the twenty-first century, ideas remain the dominion of people—and the frontier for thinking machines.

David Ferrucci's mission was to explore that frontier. Like many in his profession, Ferrucci grew up watching *Star Trek* on television. The characters on the show, humans and pointy-eared Vulcans alike, spoke to their computer as if it were one of them. No formatting was necessary, no key words, no programming language. They spoke English. The computer understood the meaning and context of the questions. It consulted vast databases and came back with an immediate answer. True, it might not produce original ideas. But it was an extravagantly well-informed shipmate. That was the computer Ferrucci wanted to build.

As he served the last drops of his wine, Ferrucci was talking about the world he was busy creating, one in which people and their machines often appeared to switch roles. He didn't know, he said, whether engineers would ever be able to "create a sentient being." But when he looked at his fellow humans through the eyes of a computer scientist, he saw pat-

programmed. He mentioned the zombielike commutes, the retreat to the same chair, the hand reaching for the TV remote, and the near-identical routines, from toothbrushing to feeding the animals. "It's more interesting," he said, "when humans delve inside themselves and say, 'Why am I doing this? And why is it relevant and important to be human?'" His machine would nudge people toward that line of inquiry. Even with an avatar for a face and a robotic voice, the *Jeopardy* machine would invite comparisons to the other two contestants on the stage. This was inevitable. And whether it won or lost on a winter evening in 2011, the computer might lead millions of spectators to reflect on the nature, and probe the potential, of their own humanity.

2.3 NANOTECHNOLOGY

2.3.1 NEUROSCIENCE, NANOTECHNOLOGY, AND ETHICS: PROMISE & PERIL

RAY KURZWEIL

Inventor, entrepreneur, author, and futurist, Ray Kurzweil is one of the leading thinkers about contemporary and future technology and its impact on society. In this reading, he argues that the pace of technological change is exponential rather than linear, as many people implicitly believe, and that the twenty-first century will witness a major technological revolution springing from the convergence of the sciences of genetics, nanotechnology, and robotics (GNR) and artificial intelligence. However, like all new technologies, GNR is a "double-edged sword" that holds great promise and great peril for human civilization. One worry, is that self-replicating nanobots will escape into the environment and cause severe and irreparable damage to the natural world. Although he takes such threats seriously, Kurzweil believes that the development and introduction of GNR is inevitable and that broad relinquishment of these technologies is not feasible, but that it is possible to develop ethical guidelines for "fine-grained" relinquishment of certain kinds of particularly dangerous GNR applications. While continuing to be optimistic about the promise of these technologies, he believes that it is possible for society to control them so as to enjoy their benefits while avoiding the dangers they pose.

FOCUS QUESTIONS

1. What are the three features of technological evolution that lead Kurzweil to conclude that the rate of technological change is exponential rather than linear?
2. Kurzweil suggests that the introduction of some kinds of nanotechnology devices is inevitable. What reasons does he adduce that support this conclusion?

Source: From the *Encyclopedia of Science, Technology, and Ethics*, Vol. 1–c, 2006, Macmillan Reference,

3. At one point, Kurzweil refers to the views of Ted Kaczynski (aka the Unabomber) and, perhaps surprisingly, seems to endorse some aspects of his view of technology. What part of Kaczynski's view does he share? What part does he disagree with?
4. Why is Kurzweil opposed to the relinquishment of broad areas of technology? What kinds of "fine-grained" relinquishment does he think are feasible?
5. Do you agree with Kurzweil that it is possible for society to enjoy the benefits of twenty-first century GNR technologies while mitigating and controlling the risks? Explain.

KEYWORDS

defensive technologies, ethical legal and social implications, genetic nanotechnology and robotics, molecular assemblers, Moore's law, nanobots, nanotechnology, paradigm shift, relinquishment, unified system

Our rapidly growing scientific and technological ability to manipulate matter and energy at ever smaller scales promises to transform virtually every sector of society, a phenomenon that presents manifest ethical responsibilities. There will be increasing overlap between nanotechnology and other technologies, such as biotechnology and artificial intelligence. And as with these previous scientific and technological transformations, we will be faced with deeply intertwined promise and peril.

THE NANO-FRONTIER

Nanoscience and nanotechnology today have been expanded to include essentially any science or technology where the key features are measured in a modest number of nanometers (under 100 by some definitions). By this standard, contemporary electronics has already passed this threshold. Eric Drexler has further developed the concept of building molecule-scale devices using molecular assemblers that would precisely guide chemical reactions by means of information. Moreover, just as technologies related to information develop at an exponential pace, generally doubling in capability and price-performance every year, so the size of technology is itself inexorably shrinking, and most of technology will be "nanotechnology" by the 2020s.

This era will bring us the ability to essentially convert software, that is, information, directly into physical products. We will be able to produce virtually any product for pennies per pound. Computers will have greater computational capacity than the human brain, and we will be completing the reverse engineering of the human brain to reveal the software design of human intelligence. We are already placing devices with narrow intelligence in our bodies for diagnostic and therapeutic purposes. With the advent of nanotechnology, we will be able to keep our bodies and brains in a healthy, optimal state indefinitely. Nanotechnology and related advanced technologies will bring us the opportunity to overcome age-old problems, including pollution, poverty, disease, and aging.

Many object to the intermingling of the so-called natural world with the products of our technology. However, the increasing intimacy of our human lives with our technology is not a new story. Human life expectancy was thirty-seven years in 1800. Most humans at that time lived lives dominated by poverty, intense labor, disease, and misfortune. We are immeasurably better off as a result of technology, but there is still a lot of suffering in the world to overcome. We have a moral imperative, therefore, to continue the pursuit of knowledge and of advanced technologies that can continue to overcome human affliction. There is also an economic imperative to continue.

Nanotechnology is advancing on hundreds of fronts. We cannot relinquish its pursuit without essentially relinquishing all of technology, which would require acts of totalitarianism inconsistent with the values of our society. Technology has always been a double-edged sword, and that is certainly true of nanotechnology. However, we will have no choice but to confront the challenge of guiding nanotechnology in a constructive

direction. Any broad attempt to relinquish nanotechnology will only push it underground, which would interfere with the benefits while actually making the dangers worse.

With the human genome project, three to five percent of the budgets were devoted to the ethical, legal, and social implications (ELSI) of the technology. A similar commitment for nanotechnology would be appropriate and constructive. Near-term applications of nanotechnology are more limited in their benefits and more benign in their potential dangers. We cannot say a priori that all nanoengineered particles are safe, nor would it be appropriate to deem them necessarily unsafe. Environmental tests thus far have not shown reasons for undue concern.

I believe that existing regulatory mechanisms are sufficient to handle near-term applications of nanotechnology. As for the long term, we need to appreciate that a myriad of nanoscale technologies are inevitable. The current examinations and dialogues on achieving the promise while ameliorating the peril are appropriate and will deserve increased attention as we get closer to realizing these revolutionary technologies.

THE NANO-BACKGROUND: MODELS OF TECHNOLOGY TRENDS

Models of technology trends show that nanotechnology and related advanced technologies are inevitable. They are deeply integrated into our society and are advancing on many diverse fronts, comprised of hundreds of small steps, each benign in itself.

Intuitive Linear and Historical Exponential Views

Although exponential trends did exist a thousand years ago, they were at that very early stage where it is so flat and so slow that it looks like no trend at all. Today, everyone expects continuous technological progress and the social repercussions that follow. But the future will nonetheless be far more surprising than most observers realize because few have internalized the fact that the rate of change itself is accelerating.

Most long-range forecasts of technical feasibility underestimate the power of future developments because they are based on the "intuitive linear" view of history rather than the "historical exponential" view. We will not experience a hundred years of progress in the twenty-first century; rather we will witness on the order of twenty thousand years of progress (at today's rate of progress). An unexamined intuition provides the impression that progress changes at the rate that we have recently experienced because an exponential curve approximates a straight line when viewed for a brief duration.

But an assessment of the history of technology shows that technological change is exponential. Indeed, we find "double" exponential growth, meaning that the rate of exponential growth is itself growing exponentially. These observations are based on a rich model of diverse technological processes.

The Law of Accelerating Returns

The ongoing acceleration of technology is the inevitable result of the "law of accelerating returns," which describes the acceleration of the pace and the exponential growth of the products of an evolutionary process, including technology, particularly information technologies.

The law of accelerating returns has three key features. First, evolution applies positive feedback as the more capable methods resulting from one stage of evolutionary progress are used to create the next stage. As a result, the rate progress of an evolutionary process increases exponentially over time, as the "returns" of that process (e.g., speed or cost-effectiveness) increase exponentially. As an evolutionary process becomes more effective, greater resources are invested in it, resulting in a second level of exponential growth (i.e., the rate of exponential growth itself grows exponentially).

A second feature is "technological paradigm shifts." A specific paradigm (a method or approach to solving a problem) provides exponential growth until the method exhausts its potential. When this happens, a paradigm shift (a fundamental change in the approach) occurs, which enables exponential growth to continue. Each paradigm follows an

"S-curve," which consists of slow growth, followed by rapid growth, followed by a leveling off as the particular paradigm matures. During this third phase in the life cycle of a paradigm, pressure builds for the next paradigm shift. The acceleration of the overall evolutionary process proceeds as a sequence of S-curves, and the overall exponential growth consists of this cascade of S-curves.

A third key feature is that the resources underlying the exponential growth of an evolutionary process are relatively unbounded. One resource is the order of the evolutionary process itself. Each stage of evolution provides more powerful tools for the next. The other required resource is the "chaos" of the environment in which the evolutionary process takes place and which provides the options for further diversity. In technological evolution, human ingenuity and the ever-changing market sustain innovation.

The evolution of life forms and technologies constantly accelerates. With the advent of a technology-creating species, the exponential pace became too fast for evolution through DNA-guided protein synthesis and moved on to human-created technology. Technology goes beyond mere tool making; it is a process of creating ever more powerful technology using the tools from the previous round of innovation. The first technological steps took tens of thousands of years. For people living in this era, there was little noticeable technological change. By 1000 C.E., progress was much faster and a paradigm shift required only a century or two. The nineteenth century saw more technological change than in the nine centuries preceding it. Then in the first twenty years of the twentieth century, we saw more advancement than in all of the nineteenth century. Now, paradigm shifts occur in only a few years. The paradigm shift rate is currently doubling every decade. So the twenty-first century will see about a thousand times greater technological change than its predecessor.

Moore's Law and Beyond

The exponential trend that has gained the greatest public recognition has become known as "Moore's

integrated circuits, noted in the mid-1970s that we could squeeze twice as many transistors on an integrated circuit every twenty-four months. Given that the electrons have less distance to travel, the circuits also run twice as fast, providing an overall quadrupling of computational power.

However, the exponential growth of computing is much broader than Moore's Law. If we plot the speed per price of forty-nine famous calculators and computers spanning the twentieth century, we note that there were four paradigms that provided exponential growth in the price-performance of computing before integrated circuits. Therefore, Moore's Law was the fifth paradigm to exponentially grow the power of computation. When Moore's Law reaches the end of its S-curve, the exponential growth will continue with three-dimensional molecular computing, constituting the sixth paradigm.

Moore's Law narrowly refers to the number of transistors on an integrated circuit of fixed size. But the most appropriate measure to track is computational speed per unit cost. This takes into account many levels of innovation in computer design. For example, there are many nascent technologies that build circuitry in three dimensions in a way that mimics the parallel organization of the human brain. One cubic inch of nanotube circuitry would be a million times more powerful than the human brain. There are more than enough new computing technologies now being researched to sustain the law of accelerating returns as applied to computation.

Specific paradigms do ultimately reach levels at which exponential growth is no longer feasible. That is why Moore's Law is an S-curve. But the growth of computation will continue exponentially. Paradigm shift, or innovation, turns the S-curve of any specific paradigm into a continuing exponential. A new paradigm takes over when the old paradigm approaches its natural limit.

Other Technologies

There are many examples of the exponential growth implied by the law of accelerating returns in technologies as varied as DNA sequencing, com-

all kinds, and even in the rapidly shrinking size of technology. Exponential growth in communications technology has been even more explosive than in computation. Miniaturization is a trend that will have profound implications for the twenty-first century. The salient implementation sizes of technologies, both electronic and mechanical, are shrinking at a double-exponential rate.

The future nanotechnology age will result not from the exponential explosion of computation alone, but rather from the synergies that will result from intertwined technological revolutions. Every point on the exponential growth curves represents an intense human drama of innovation and competition. It is remarkable that these chaotic processes result in such smooth and predictable exponential trends.

EXAMPLES OF TRUE NANOSCIENCE AND NANOTECHNOLOGY

Ubiquitous nanoscience and nanotechnology is two to three decades away. One forthcoming achievement will be "nanobots," small robots the size of human blood cells that can travel inside the human bloodstream. There have already been successful animal experiments using this concept.

In addition to human brain reverse engineering, these nanobots will be able to perform a broad variety of diagnostic and therapeutic functions inside the human body. Robert Freitas, for example, has designed robotic replacements for human blood cells that perform thousands of times more effectively than their biological counterparts. His "respirocytes" (robotic red blood cells) could allow one to sprint for fifteen minutes without taking a breath. His robotic macrophages will be far more effective than our white blood cells at combating pathogens. His DNA repair robot would be able to repair DNA transcription errors, and even implement needed DNA changes. Although Freitas' conceptual designs are two or three decades away, there has already been progress on bloodstream-based devices.

Nanobot technology has profound military applications, and any expectation that such uses will be relinquished is highly unrealistic. Already, the U.S. Department of Defense (DOD) is developing "smart dust," or tiny robots to be used for surveillance. Billions of invisible spies could monitor every square inch of enemy territory and carry out missions to destroy enemy targets. The only way for an enemy to counteract such a force is with their own nanotechnology. Nanotechnology-based weapons will obsolete weapons of larger size.

In addition, nanobots will be able to expand our experiences and our capabilities. Nanobot technology will provide fully immersive virtual reality by taking up positions in close proximity to every interneuronal connection related to the senses. If we want to enter virtual reality, the nanobots suppress all of the inputs coming from the real senses, and replace them with the signals that would be appropriate for the virtual environment.

Scientists at the Max Planck Institute have developed "neuron transistors" that can detect the firing of a nearby neuron, or alternatively, can cause a nearby neuron to fire, or suppress it from firing. This amounts to two-way communication between neurons and the electronic-based neuron transistors. The scientists demonstrated their invention by controlling the movement of a living leech from their computer.

The Internet will provide many virtual environments to explore. We will be able to "go" to these virtual environments and meet others there, both real and simulated people. Of course, ultimately there will not be a clear distinction between the two. By 2030, going to a web site will mean entering a full-immersion virtual-reality environment, encompassing all of the senses and triggering the neurological correlates of emotions and sexual experiences.

"Experience beamers" circa 2030 will beam a person's entire flow of sensory experiences and emotions. We'll be able to go to a web site and experience other people's lives. Full-immersion visual-auditory environments will be available by 2010, with images written directly onto our retinas by our eyeglasses and contact lenses. The electronics will be embedded in our glasses and woven into our clothing, so computers as distinct objects will disappear.

The most significant implication of nanotechnology and related advanced technologies of the twenty-first century will be the merger of biological and nonbiological intelligence. Nonbiological intelligence is growing at a double-exponential rate and will vastly exceed biological intelligence well before the middle of this century. However, in my view, this nonbiological intelligence should still be considered human, as it is fully derivative of the human-machine civilization.

Our brains are relatively fixed in design, but brain implants based on massively distributed intelligent nanobots will ultimately expand our memories a trillion fold and improve all of our cognitive abilities. Since the nanobots are communicating with each other over a wireless network, they can create any set of new neural connections, break existing connections, create new hybrid biological-nonbiological networks, and add new nonbiological networks.

Using nanobots as brain extenders is a significant improvement over surgically installed neural implants. Nanobots will be introduced without surgery and can be directed to leave, so the process is easily reversible. They can change their configuration and alter their software. Perhaps most importantly, they are massively distributed and can take up billions or trillions of positions throughout the brain, whereas a surgically introduced neural implant can only be placed in a few locations.

THE ECONOMIC IMPERATIVES OF THE LAW OF ACCELERATING RETURNS

The economic imperative of a competitive marketplace is driving science and technology forward and fueling the law of accelerating returns, which, in turn, is transforming economic relationships. We are moving toward nanoscale, more intelligent machines as the result of many small advances, each with their own particular economic justification.

There is a vital economic imperative to create smaller and more intelligent technology. Machines that can more precisely carry out their missions have enormous value. There are tens of thousands of projects that are advancing the various aspects of the law of accelerating returns in diverse incremental ways. Regardless of near-term business cycles, the support for "high tech" in the business community has grown enormously. We would have to repeal capitalism and every visage of economic competition to stop this progression.

The economy has been growing exponentially throughout this century. Even the Great Depression of the 1930s represented only a minor blip compared to the underlying pattern of growth. Recessions, including the Depression, represent only temporary deviations from the underlying curve. Statistics in fact greatly understate productivity growth (economic output per worker), which has also been exponential.

Inflationary factors are offset by the double-exponential trends in the price-performance of all information-based technologies, which deeply affect all industries. We are also undergoing massive disintermediation in the channels of distribution through the Internet and other new communication technologies and escalating efficiencies in operations and administration. Current economic policy is based on outdated theories that do not adequately model the size of technology, bandwidth, megabytes, intellectual property, knowledge, and other increasingly vital constituents that are driving the economy.

Cycles of recession will not disappear immediately. However, the rapid dissemination of information, sophisticated forms of online procurement, and increasingly transparent markets in all industries have diminished the impact of these cycles. The underlying long-term growth rate will continue at a double-exponential rate. The rate of paradigm shift is not noticeably affected by the minor deviations caused by economic cycles. The overall growth of the economy reflects completely new forms of wealth and value that did not previously exist: nanoparticle-based materials, genetic information, intellectual property, communication portals, web sites, bandwidth, software, data bases, and many other new technology-based categories.

Another implication of the law of accelerating returns is exponential growth in human knowledge, including intellectual property, education,

and learning. Over the course of the long twentieth century we increased investment in K-12 education by a factor of ten. We have a one hundred fold increase in the number of college students. Automation has been eliminating jobs at the bottom of the skill ladder while creating new and better paying jobs at the top. So, the ladder has been moving up, and we have been exponentially increasing investments in education at all levels.

PROMISE AND PERIL

Science and technology have always been double-edged swords, bringing us longer and healthier life spans, freedom from physical and mental drudgery, and many new creative possibilities, while at the same time introducing new and salient dangers. We will need to adopt strategies to encourage the benefits while ameliorating the risks. Relinquishing broad areas of technology, as some critics have proposed, is not feasible, and attempts to do so will only drive technology development underground, which will exacerbate the dangers.

As technology accelerates toward the full realization of biotechnology, nanotechnology and "strong" AI (artificial intelligence at or above human levels), we will see the same intertwined potentials: a feast of creativity resulting from greater human intelligence combined with many new dangers. Nanobot technology requires billions or trillions of such intelligent devices to be useful. The most cost-effective way to scale up to such levels is through self-replication. A defect in the mechanism curtailing nanobot self-replication could be disastrous. There are steps available now to mitigate this risk, but we cannot have complete assurance in any strategy that we devise today.

Other primary concerns include "Who is controlling the nanobots?" and "Who are the nanobots talking to?" Organizations or individuals could put undetectable nanobots in water or food supplies. These "spies" could monitor and even control thoughts and actions. Existing nanobots could be influenced through software viruses and other software "hacking" techniques. My own expectation is that the creative and constructive applications of this technology will dominate, as they do today. But we need to invest more heavily in developing specific defensive technologies.

There are usually three stages in examining the impact of future technology: awe at its potential to overcome problems; then a sense of dread at a new set of dangers; followed by the realization that the only viable and responsible path is to set a careful course that can realize the promise while managing the peril.

Bill Joy, cofounder of Sun Microsystems, has warned of the impending dangers from the emergence of self-replicating technologies in the fields of genetics, nanotechnology, and robotics, or "GNR." His concerns include genetically altered designer pathogens, self-replicating entities created through nanotechnology, and robots whose intelligence will rival and ultimately exceed our own. Who's to say we will be able to count on such robots to remain friendly to humans? Although I am often cast as the technology optimist who counters Joy's pessimism, I do share his concerns regarding self-replicating technologies. Many people have interpreted Joy's article as an advocacy of broad relinquishment, not of all technology, but of the "dangerous ones" like nanotechnology. Joy, who is now working as a venture capitalist with the legendary Silicon Valley firm of Kleiner, Perkins, Caufield & Byers investing in technologies such as nanotechnology applied to renewable energy and other natural resources, says that broad relinquishment is a misinterpretation of his position and was never his intent. He has recently said that the emphasis should be to "limit development of the technologies that are too dangerous," not on complete prohibition. He suggests, for example, a prohibition against self-replicating nanotechnology, which is similar to the guidelines advocated by the Foresight Institute.

Others, such as Bill McKibben, the environmentalist who was one of the first to warn against global warming, have advocated relinquishment of broad areas such as biotechnology and nanotechnology, or even of all technology. However, relinquishing broad fields would be impossible to achieve without essentially relinquishing all technical development.

There are real dangers associated with new self-replicating technologies. But technological advances, such as antibiotics and improved sanitation, have freed us from the prevalence of such plagues in the past. We may romanticize the past, but until fairly recently, most of humanity lived extremely fragile lives. Many people still live in this precarious way, which is one reason to continue technological progress and the economic enhancement that accompanies it. Should we tell the millions of people afflicted with devastating conditions that we are canceling the development of all bioengineered treatments because there is a risk that these same technologies may someday be used for malevolent purposes? Most people would agree that such broad-based relinquishment is not the answer.

The Relinquishment Issue

Relinquishment at the right level is part of a responsible and constructive response to these genuine perils. The issue, however, is: At what level are we to relinquish technology? Ted Kaczynski (the Unabomber) would have us renounce all of it. This is neither desirable nor feasible. McKibben takes the position that many people now have enough wealth and technological capability and should not pursue more. This ignores the suffering that remains in the human world, which continued technological progress could alleviate.

Another level would be to forego certain fields (such as nanotechnology) that might be regarded as too dangerous. But such sweeping strokes of relinquishment are untenable. Nanotechnology is the inevitable result of the persistent trend toward miniaturization that pervades all of technology. It is not a single centralized effort, but is being pursued by a myriad of projects with many goals.

Kaczynski argued that modern industrial society cannot be reformed because technology is a unified system in which all parts are dependent on one another. It is not possible to get rid of the "bad" parts of technology and retain only the "good" parts. He cited modern medicine as an example, arguing that progress depends on several scientific fields and advancements in high-tech equipment. Kaczynski was correct on the deeply entangled nature of the benefits and risks, but his overall assessment of the relative balance between the two was way off. Joy and I both believe that technology will and should progress, and that we need to be actively concerned with the dark side. Our dialogue concerns the granularity of relinquishment that is feasible and desirable. Abandonment of broad areas of technology will only push them underground where development would continue unimpeded by ethics and regulation. In such a situation, it would be the less-stable, less-responsible practitioners who would have all the expertise.

One example of relinquishment at the right level is the proposed ethical guideline by the Foresight Institute that nanotechnologists agree to relinquish the development of physical entities that can self-replicate in a natural environment. Another is a ban on self-replicating physical entities that contain their own codes for self-replication. Such entities should be designed to obtain codes from a centralized secure server, which would guard against undesirable replication. This "broadcast architecture" is impossible in the biological world, which represents one way in which nanotechnology can be made safer than biotechnology. Such "finegrained" relinquishment should be linked to professional ethical guidelines, oversight by regulatory bodies, the development of technology-specific "immune" responses, as well as computer-assisted surveillance by law enforcement agencies. Balancing privacy rights with security will be one of many challenges raised by some new nanotechnologies.

Computer viruses serve as a reassuring test case in our ability to regulate nonbiological self-replication. At first, concerns were voiced that as they became more sophisticated, software pathogens had the potential to destroy computer networks. Yet the "immune system" that has evolved in response to this challenge has been largely effective. Although self-replicating software entities do cause damage from time to time, no one would suggest we do away with computers and the Internet because of software viruses. This success is in a highly productive industry in

which there is no regulation, and no certification for practitioners.

Defensive Technologies and The Impact of Regulation

Arguments such as McKibben's for relinquishment have been influential because they paint a picture of future dangers as if they were released into an unprepared world. But the sophistication and power of our defensive technologies and knowledge will grow along with the dangers. When we have "gray goo" (unrestrained nanobot replication), we will also have "blue goo" ("police nanobots"). We cannot say with assurance that we will successfully avoid all misuse. We have been able to largely control harmful software virus replication because the requisite knowledge is widely available to responsible practitioners. Attempts to restrict this knowledge would have created a far less stable situation.

The present challenge is self-replicating biotechnology. By reprogramming the information processes that lead to and encourage disease and aging, we will have the ability to overcome these afflictions. However, the same knowledge can also empower a terrorist to create a bioengineered pathogen.

Unlike biotechnology, the software industry is almost completely unregulated. Although bio-terrorists do not need to put their "innovations" through the FDA, scientists developing defensive technologies are required to follow regulations that slow innovation. It is impossible under existing regulations and ethical standards to test defenses to bioterrorist agents on humans. Animal models and simulations will be necessary in lieu of infeasible human trials, but we will need to go beyond these steps to accelerate the development of defensive technologies.

We need to create ethical and legal standards and defensive technologies. It is quite clearly a race. In the software field the defensive technologies have remained ahead of the offensive ones. With extensive regulation in the medical field slowing down innovation, this may not happen with biotechnology.

There is a legitimate need to make biomedical research as safe as possible, but our balancing of risks is skewed. The millions of people who need biotechnology advances seem to carry little political weight against a few well-publicized casualties from the inevitable risks of progress. This equation will become even starker with the emerging dangers of bioengineered pathogens. We need a change in public attitude in terms of tolerance for necessary risk.

Hastening defensive technologies is vital to our security. We need to streamline regulatory procedures to achieve this. However, we also need to greatly increase our investment explicitly in defensive technologies. In the biotechnology field, this means the rapid development of antiviral medications.

The comparable situation will exist for nanotechnology once replication of nanoengineered entities has been achieved. We will soon need to invest in defensive technologies, including the creation of a nanotechnology-based immune system. Such an immune system may itself become a danger, but no one would argue that humans would be better off without an immune system because of the possibility of autoimmune diseases. The development of a technological immune system for nanotechnology will happen even without explicit efforts to create one.

It is premature to develop specific defensive nanotechnologies as long as we have only a general idea of the threat. However, there is a dialogue on this issue, and expanded investment in these efforts should be encouraged. The Foresight Institute, for example, has devised a set of ethical standards and strategies for assuring the development of safe nanotechnology. They are likely to be effective with regard to preventing accidental release of dangerous self-replicating nanotechnology entities. But the intentional design and release of such entities is more challenging.

CONCLUSION

Protection is not impossible, but we need to realize that any level of protection will only work to a certain level of sophistication. We will need to continue to advance the defensive technologies

and keep them ahead of the destructive technologies. The challenge of self-replication in nanotechnology impels us to continue the type of study that the Foresight Institute has initiated. With the human genome project, three to five percent of the budget was devoted to the ethical, legal, and social implications (ELSI) of the technology. A similar commitment for nanotechnology would be appropriate and constructive. Science and technology will remain double-edged swords, and the story of the twenty-first century has not yet been written. We have no choice but to work hard to apply these quickening technologies to advance our human values, despite what often appears to be a lack of consensus on what those values should be.

BIBLIOGRAPHY

Drexler, Erik K. (1986). *Engines of Creation*. Garden City, NY: Doubleday.

Drexler, Erik K. (1992). *Nanosystems: Molecular Machinery, Manufacturing, and Computation*. New York: Wiley.

Freitas, Robert, Jr. (1999). *Nanomedicine*, Vol. 1: *Basic Capabilities*. Georgetown, TX: Landes Bioscience.

Joy, Bill. (2000). "Why the Future Doesn't Need Us." *Wired* 262.

Kurzweil, Ray. (1999). *The Age of Spiritual Machines When Computers Exceed Human Intelligence*. New York: Viking.

Kurzweil, Ray. (2005). *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking.

McKibben, Bill. (2003). *Enough: Staying Human in an Engineered Age*. New York: Henry Holt.

2.3.2 EMERGING NANOMATERIAL GOVERNANCE SYSTEMS: THE STATE OF PLAY

LYNN L. BERGESON

As with many fields of technology, there are strong financial incentives for companies to rapidly commercialize nanotechnology products. According to a study by a nanotechnology research and advocacy firm cited by the author of this selection, nanotechnology products will constitute 15 percent of global manufacturing output, amounting to $2.6 trillion, by 2014. But despite this rapid industrial growth, there is a lack of reliable research on the expected environmental, health, and safety impacts of the widespread diffusion of nanotechnology products. The key challenge to be addressed at the present is how society can benefit from nanotechnology products while at the same time effectively manage the unknown and uncertain risks associated with exposure to nanoscale materials.

This selection reviews several types of responses to this challenge. Governmental regulatory regimes through agencies such as the U.S. Environmental Protection Agency, the National Toxicology Program, and other agencies of the federal government have looked at the issue through the lens of existing laws. But, as the author notes, no new laws that specifically address the environmental, health, and safety risks associated with nanoscale products have been passed. There have also been attempts to address this challenge at the international level, through a working group of the Organization of Economic, Co-operation and Development (OECD), but this initiative

Source: Bergeson, L. L. "Emerging Nanomaterial Governance Systems: The State of Play." *Molecular*

has also not led to any mandatory regulations on companies that are producing nanoproducts. In the face of this regulatory gap, several cities and states have attempted to develop laws that may point the way to broader legislative responses. However, for the time being, the most promising initiatives to govern the development and deployment of products containing nanoscale materials appear to be voluntary codes developed jointly by industry and nongovernmental organizations (NGOs), several of which are discussed in this selection. These codes, like other voluntary initiatives, are not legal requirements that all companies must adhere to, so they leave open the possibility that some firms may not comply with their responsibilities and go ahead with the commercialization of nanotechnology products without appropriate risk management strategies in place.

FOCUS QUESTIONS

1. Discuss why nanotechnology is seen as having a variety of beneficial aspects as well as posing potential risks. What are some of the unique aspects of this technology that makes it a particular health threat?
2. Discuss why you feel it is important that various U.S. and/or international agencies are concerned with the possible EHS effects of nanoscale materials or whether you believe the private sector is prepared to provide adequate protection to the public on their own.
3. Who are some of the key stakeholder groups whose safety might be impacted by this new technology? What mechanisms are in place today and what new measures might have to be taken to maintain the level of protection required?
4. How do the financial incentives associated with advances in nanotechnology and the globalized world in which we live pose unusual safety risks from this technology compared to earlier advances in other areas?

KEYWORDS

carbon nanotubes (CNTs), Environmental, Health, and Safety (EHS), GoodNanoGuide, insurance, Nano Risk Framework, Nanoscale Materials Stewardship Program (NMSP), nanosilver products, product liability, risk management.

NANOTECHNOLOGY is the "understanding and control of matter at dimensions of roughly 1 to 100 nanometers, where unique phenomena enable novel applications," according to the National Nanotechnology Initiative.[1] Nanotechnology is viewed broadly as many technologies over time are expected to generate many new products and applications. Lux Research, Inc., the New York-based nanotechnology research and advocacy firm, predicts that by 2014, products that incorporate nanotechnology will constitute 15% of global manufacturing output and will total $2.6 trillion.[2]

Governments globally are focusing on nanotechnology for reasons beyond economic advancement. Because there is a lack of complete understanding regarding the environmental, health, and safety (EHS) effects of exposure to engineered nanoscale materials, government bodies and industry are considering how best to address EHS issues while fostering the commercialization of nanoscale materials. It is generally believed that sufficient information exists about the potential toxicity of nanoscale materials to suggest a need for caution. The small size of certain nanoparticles facilitates their uptake into cells and their movement through the body more readily than is the case with their macrocounterparts.[3]*

* Toxicology studies of certain ultrafine particles demonstrate that smaller particles have the potential to induce oxidative stress and inflammation in the respiratory tract and cardiovascular systems.

In addition to size, however, other factors contribute to a general sense of uncertainty as to the biologic and environmental implications of exposure to engineered nanoscale materials. The complexity of engineered nanomaterials means that their impact will depend on more than chemistry alone. Size, shape, surface chemistry, and surface coating, for example, can all influence how these materials behave biologically and in the environment. In some instances, the fact that nanoscale materials can have unusual properties—properties that do not conform to "conventional" physics and chemistry—may increase their commercial value and their potential for risk.

Many federal agencies are engaged in the active review of nanotechnology applications and the EHS implications of nanotechnologies. These include the Environmental Protection Agency (EPA),[4] the Food and Drug Administration (FDA),[5] the Occupational Safety and Health Administration (OSHA),[6] the National Institute for Occupational Safety and Health (NIOSH),[7] the National Toxicology Program,[8] and the Department of Defense,[9] among other federal agencies and departments.

Regulatory and health agencies globally are similarly engaged. The EPA is most prominently involved in reviewing the EHS implications of nanotechnologies and in identifying and funding research initiatives regarding the beneficial environmental applications of nanotechnologies in a wide variety of areas.[10]

The products of nanotechnology are diverse and growing exponentially. According to the National Nanotechnology Initiative, nanoparticles and nanoscale materials are used in many industries, including electronics, pharmaceuticals, chemicals, energy, among others. Reportedly, areas producing the greatest revenue for nanoparticles are chemical and mechanical polishing, magnetic recording tapes, sunscreens, automotive catalyst supports, biolabeling, and electroconductive coatings and optical fibers. According to one source, over 1,000 manufacturer-identified nanotechnology-enabled products are currently on the market, including paints, sporting goods, cosmetics, stain-resistant clothing, electronics, and surface coatings, among other applications.[11]

The uncertainty of nanotechnology, its health and safety implications, its commercial applications, the public's receptivity to nanotechnology, and the global governance mechanisms that may be applied to manage the risks potentially arising from nanoscale materials do not play well to an increasingly cautious public or the investment and insurance industries, whose support for the successful commercialization of this technology is essential. The hope is that prudent business practices will prevail, the promise of nanotechnology will materialize, and the science will unfold in ways that support its continued commercialization. In the interim, business people will be challenged to address increasingly complex legal, commercial, regulatory, and public relations issues, as summarized below.

POTENTIAL LIABILITIES AND GOVERNANCE SYSTEMS

Nanotechnology companies face a range of potential liabilities. Assessing and quantifying these potential liabilities are difficult to do given the relative embryonic stage of the field, the lack of well-defined legal precedent and regulatory governance standards, evolving industry standards and practices, and the fast-developing science. Several areas of potential liability warrant note here.

Product Liability

The specter of product liability casts a shadow over nanotechnology innovations. The "Magic Nano" product became the poster child for nanotechnology product liability concerns following its recall and a relentless stream of bad publicity. Magic Nano was first sold in supermarkets and discount stores in March 2006 in Germany. The product was a household glass and ceramic tile sealant available in aerosol cans. According to the German Federal Institute for Risk Assessment,[12] between March 27 and March 30, 2006, almost 100 people who used the product claimed to experience breathing problems requiring hospital treatment. It was ultimately determined that the product contained no nanoparticles. That the product contained no

nanoscale materials, and decidedly no magic, is irrelevant in the court of public opinion, however, and the incident was widely regarded as a wake-up call to industry as to the mischief "nanotechnology" can inspire. Strict liability, negligence, breach of warranty, and claims made in connection with these legal principles (eg, design defect, manufacturer defect, and failure to warn) arguably could be asserted against manufacturers or suppliers of nano-enabled consumer products, given the right facts. Claims for damages could include personal injury, medical monitoring, fear of future injury, deceptive trade practices (inviting treble damages), and punitive damages.

Although no reported litigation appears to have been brought as of mid-2010 involving nanotechnology or a nano-enabled product, many believe that this could change in the not too distant future.[13] The full range of protective measures to minimize product liability must be considered, including contractual protections with upstream and downstream suppliers; implementation of best management practices; contractual representations and warranties; indemnification agreements; and appropriate and legally sound warnings, labeling, and related disclosure strategies. Additionally, businesses need to track postsale consumer product complaint and incident reports and respond to any quickly and thoroughly.

Businesses also must respond to the relentless pressure to communicate with diverse stakeholders about the technology they employ—stakeholders including workers, community residents, downstream customers, distributors, and others—each of which can be expected to seek and need to know what a nano-enabled product contains, the health implications from exposure to the product and its nano components, the end-of-life consequences of nano-enabled products, and related product composition and safety information. Stakeholders have come to expect even greater information given demanding right-to-know laws and the expectations they invite.

In addition to these right-to-know disclosure pressures, increasingly demanding shareholder rights activists can be expected to continue to insist on enhanced transparency and product disclosure with regard to a wide array of issues, particularly environmental and product toxicity issues. On December 11, 2008, for example, a coalition of activists wrote then President-Elect Obama asking that investors be afforded a right to propose and vote on resolutions directing a company to evaluate how specific risks may affect the company's business. Among the types of risks the coalition specifically identified were "climate change and product toxicity."[14]

Similarly, the advocacy group As You Sow Foundation's Corporate Social Responsibility Program noted in early 2009 that product safety shareholder resolutions urging companies to disclose the presence of nanomaterials in food and personal care products have greatly increased in the recent past.[15] According to the Investor Environmental Health Network, shareholders filed or refiled 62 resolutions at 39 companies for consideration at shareholder meetings between 2008 and 2009.[16,17] Many of these specifically requested information on nanomaterials.

Although the Obama administration has yet to take up this type of disclosure, the EPA has made greater transparency a core agency priority. On December 23, 2009, for example, the EPA requested comment on options for disclosing inert ingredients in pesticides[18†] and ideas for greater disclosure of inert ingredient information. Inert ingredients are part of the end-use product formulation and are not active ingredients.

As another example, on January 21, 2010, the EPA announced a new practice concerning confidential business information (CBI) claims for substantial risk information submitted to the EPA under Toxic Substances Control Act (TSCA) Section 8(e). According to the EPA's announcement, if a chemical substance is listed on the public portion of the TSCA inventory, the EPA expects that a company submitting a health and safety study to the EPA under TSCA Section 8(e) for that chemical substance will not

† The EPA is considering two general types of approaches to increasing public availability of inert ingredient identities. One would mandate disclosure only of potentially hazardous ingredients, and the other would promote or mandate public availability of most or all inert ingredient identities, regardless of hazard.

claim the chemical identity as confidential. According to the EPA, "[t]his action is part of a broader effort to increase transparency and provide more valuable information to the public by identifying programs where non-CBI may have been claimed and treated as CBI in the past."[19]

More recently, on May 27, 2010, the EPA announced that it will begin a "general practice" of reviewing confidentiality claims for chemical identities included in health and safety studies and in data from health and safety studies, submitted under the TSCA in accordance with EPA regulations at 40 C.F.R. Part 2, Subpart B.[20] According to the EPA, TSCA Section 14(b) does not extend confidential treatment to health and safety studies or data from health and safety studies that, if made public, would not disclose confidential processes used in the manufacturing or processing of a chemical substance or mixture or, in the case of a mixture, the release of data disclosing the portion of the mixture composed of any of the chemical substances in the mixture. Where a chemical identity does not explicitly contain confidential process information or reveal portions of a mixture, the EPA stated that it expects to find that the information would clearly not be entitled to confidential treatment.

Insurance

The uncertainties noted above help explain the formidable challenges that insurance coverage presents both to carriers in providing coverage and to businesses in obtaining it. The key issue in insuring against liabilities where nanotechnology is involved is the relative lack of certainty regarding potential EHS risks. Swiss Re's often-cited report, *Nanotechnology: Small Matter, Many Unknowns,* bluntly states that, in the case of nanotechnology, uncertainties prevail because "neither the probability nor the extent of the potential losses are precisely calculable," and businesses can expect to see creative approaches by the insurance industry as it tries to address a growing need among an expanding client base that relies on nanotechnology.[21] This may mean that some risks will be explicitly excluded from coverage because some essential element in the calculus is too speculative.

In April 2010, Lexington Insurance introduced what appears to be a first-of-its-kind product called LexNano-Shield. The product is "integrated liability insurance" that extends explicit coverage for nanotechnology exposures arising out of general liability, product liability, product pollution legal liability, and product recall liability. LexNanoShield claims to reimburse first-party product recall expenses incurred if a product containing nanoparticles or nanomaterials is recalled from the market for safety reasons. According to the website, LexNanoShield provides a broad array of services.[22†]

Other options that carriers may consider pursuing to address nanotechnology's undefined potential risks include capping coverage at one end or the other (insurers either will pay claims only up to a specified ceiling or will pay the excess only after the insured shoulders a large deductible); other types of risk-pooling agreements; and limitations and exclusions narrowly defined, in the form of bond sales or purchases, among other options. This is an emerging area, and no hard and fast rules yet apply.

Evolving Governance Mechanisms

Traditional governance mechanisms in the United States such as law enactments and/or standard Administrative Procedure Act notice and comment rule makings are thought by some to be challenging and possibly ill-suited tools for addressing potential EHS risks posed by the lightning-fast pace of evolving nanotechnology. As a potential matter, no laws specific to nanoscale materials have been enacted. Most government agencies in the United States and worldwide are of the view

† These include up to 2 hours of confidential consultation with a national law firm with expertise in nanotechnology to assist with appropriate risk-allocating contract language, product labeling, product warranties, development and/or review of a nanotechnology risk management program, and/or participation in the EPA Nanoscale Materials Stewardship Program; up to 2 hours of outside technical consulting services in the areas of toxicology or EHS to help with developing, implementing, or assessing a nanotechnology risk management program; and on-site loss control consultations specific to mitigating hazards of nanomaterials used in the manufacturing process or included in the product.

that they lack sufficient data and information to make informed judgments on the specific types of potential hazards and risks posed by some nanoscale materials and hence are reluctant to enact laws imposing restrictions that may or may not be warranted. In the interim, the EPA—and regulatory bodies globally—is pursuing more innovative governance strategies to assist in addressing the potential risks associated with nanotechnology. A few of these measures are described below.

US GOVERNMENT-SPONSORED VOLUNTARY INITIATIVES

In the absence of data on nanoscale materials, the EPA under the Bush administration created a voluntary disclosure process called the Nanoscale Materials Stewardship Program (NMSP), managed by the Office of Pollution Prevention and Toxics. In preparation for the program, the EPA developed a draft NMSP concept paper that outlined its "initial thinking on the design and development" of the NMSP and explained, among other matters, that the program would consist of two parts: a "Basic Program" and an "In-Depth Program."[23] The EPA solicited comment on the concept paper.

After reviewing comments on the NMSP documents, the EPA formally launched the NMSP on January 28, 2008. Under the first part of the NMSP, the Basic Program, participants were invited voluntarily to report available information on the engineered nanoscale materials they manufacture, import, process, or use. Under the In-Depth Program, participants voluntarily develop data over a longer period of time, alone or in consortia, for a particular nanoscale material.

On January 12, 2009, the EPA released its interim report on the NMSP.[24] The EPA stated that, based on the current interim results, "the NMSP can be considered successful." The EPA noted that a number of EHS data gaps still exist, however, and it "is considering how to best use testing and information gathering authorities under [the TSCA] to help address those gaps." The EPA stated then, and has reaffirmed its view,[25] that it will continue to review new chemical nanoscale materials submitted under TSCA Sections 5(a) and 5(h)(4) and apply, as appropriate, testing requirements and exposure controls under Section 5(e) and significant new use rules (SNURs) under Section 5(a)(2). The EPA's NMSP concluded in December 2009.

The NMSP generally has been viewed as a failure. EPA Assistant Administrator for Toxics Steve Owens called the voluntary program "less than a resounding success" on September 10, 2008.[26] The NMSP, on the whole, did not provide the information that the EPA expressed interest in obtaining. As a result, the EPA has opted to avail itself of more traditional governance tools, most notably rule makings under the TSCA and enforcement of existing laws, which are discussed infra.

The EPA also works extensively with international organizations on voluntary initiatives to understand and address the environmental applications and implications of nanotechnology. The Organization for Economic Cooperation and Development (OECD) established the Working Party on Manufactured Nanomaterials (WPMN), and the EPA is actively participating in the WPMN and contributes to all of these projects. Of particular relevance is the project on safety testing of a representative set of manufactured nanomaterials. The WPMN has identified a representative list of manufactured nanoscale materials for EHS testing, including fullerenes (C_{60}), single- and multiwalled carbon nanotubes (CNTs), silver nanoparticles, iron nanoparticles, carbon black, titanium dioxide, aluminum oxide, cerium oxide, zinc oxide, silicon dioxide, polystyrene, dendrimers, and nanoclays. The WPMN has also published a list of testing end points in the following areas: nanomaterial information and identification, physical and chemical properties, material characterization, environmental fate, environmental toxicology, mammalian toxicology, and material safety.

The EPA's Office of Pesticide Programs (OPP) also is active in nanotechnology. The OPP formed an intraoffice workgroup of 20 members several years ago to develop a regulatory framework and to assist in the examination of hazard, exposure, policy, regulatory, and international issues

arising in connection with nanoscale materials used as pesticides, including antimicrobial pesticides.[§] The OPP's Nanotechnology Workgroup has the job of addressing issues pertinent to nanopesticides. The OPP reportedly began receiving inquiries about registering nanoscale pesticides in 2006, particularly involving antimicrobial uses, and has received at least four formal applications to register a nanoscale pesticide—all antimicrobial products, as discussed below.

The OPP is also coordinating with the OECD Working Party on Pesticides and Task Force on Biocides to develop a survey to gather basic information from OECD member countries on their respective involvement with pesticides and biocides and nanotechnology and to identify the various OECD member countries' regulatory approaches to nanotechnology-related pesticide and biocide issues.

FEDERAL EPA REGULATORY INITIATIVES

The EPA's regulatory initiatives have picked up steam significantly over the past several years. To assist manufacturers in understanding the regulatory status of chemical nanoscale materials, the EPA prepared a policy statement dated January 2008, *TSCA Inventory Status of Nanoscale Substances — General Approach*.[27] On October 31, 2008, the EPA published an important notice outlining the TSCA requirements potentially applicable to nanoscale materials and advised manufacturers of CNTs of its position that CNTs must be listed on the TSCA inventory.[28] After March 1, 2009, CNTs that are manufactured for commercial purposes and are not listed on the TSCA inventory or are otherwise exempt may be the subject of compliance-monitoring efforts.

The EPA has also received, reviewed, and taken action on many new chemical notices under TSCA Section 5 pertinent to nanoscale materials reflecting enhanced commercial activity with respect to nanoscale materials. On May 7, 2008, the EPA issued a *Federal Register* notice announcing receipt of a premanufacture notification (PMN) to manufacture single-walled CNTs.[29] On November 5, 2008, the EPA issued a final SNUR for 56 substances, 2 of which included nanoscale substances.[30] On January 5, 2009, the EPA announced receipt of several premanufacture notifications concerning multiwalled CNTs.[31] On June 24, 2009, the EPA issued a direct final rule promulgating SNURs from multi- and single-walled CNTs.[32] In light of the adverse public comment, on August 21, 2009, the EPA withdrew the direct final rule[33] and on February 3, 2010, proposed the SNUR.[34]

On the FIFRA side of the EPA's Office of Pollution Prevention and Toxics, the EPA announced in April 2010 that sometime in 2010 it intends to adopt a policy that would require any pesticide registrant that is aware that some constituent of a registered pesticide product is nano-sized (ie, has particles or structures with a diameter less than 100 nm) to submit the information to the EPA pursuant to FIFRA Section 6(a)(2).[35] EPA regulations generally limit the obligation of a registrant to report information pursuant to FIFRA Section 6(a)(2) to information that concerns "adverse effects," so this expansion of EPA reporting requirements appears to be based on a premise that EPA regards the mere presence of any nanoscale materials to be "adverse."

From November 3 to 5, 2009, the FIFRA Scientific Advisory Panel (SAP) met to consider for the first time scientific issues related to the assessment of potential hazard and exposure associated with nanosilver and other nano-metal pesticide products. The decision to convene a SAP may have been prompted in part by a May 1, 2008, petition for rule making filed by the International Center for Technology Assessment (ICTA) requesting that the EPA regulate products containing nanoscale silver as pesticides. The petition expressed concern about the potential human health and environmental risks of nanoscale materials and requested that the EPA

[§] The Workgroup is divided into four subgroups. It includes individuals with expertise in chemistry, environmental law and policy, toxicology, exposure and risk assessment, and other areas. The Workgroup has focused particularly on the potential exposure and hazards of nanoscale pesticides and how these concerns may or may not be addressed by traditional testing paradigms and risk assessment approaches.

classify nanoscale silver as a pesticide and require registration. The EPA requested comment on the ICTA petition on November 19, 2008.[36] The OPP noted that it intended to issue a response to the petition after reviewing any public comment and coordinating with the EPA's Office of Enforcement and Compliance Assurance. More information about pesticide nanotechnology issues is available on the EPA website.[37]

Since the ICTA filed the petition, the EPA has been asked to consider at least four applications seeking registration of products containing nanosilver-based active ingredients. The nanosilver products would take the form of textile additives, polymers, coatings, and/or plastics and would be used to protect a treated product from microorganisms or to impart antimicrobial activity to a treated material. These products would be used in the same manner as some of the currently registered silver products, including those used as material preservatives and antimicrobial pesticides.

In its background paper to the SAP meeting,[38] the EPA states that the current state of the science does not contain sufficient information to determine definitively whether (and, if so, to what extent) various forms of nanosilver particles may cause toxic effects beyond those attributable to the release of silver ions. In light of this, the threshold question before the SAP relates to whether the EPA can make its safety finding under the FIFRA that a pesticide product will not cause unreasonable adverse effects on the environment with respect to the four pending applications.

According to the EPA, the registration applicants assert that the mode of action for nanosilver is the same as for silver in that the release of silver ions is the source of antimicrobial activity. Because the pesticidal mode of action of nanosilver is the same as for conventionally sized silver, the potential hazards to human health and the environment resulting from the use of nanosilver as a pesticide will, therefore, be the same as from the use of silver. The EPA likened the registrants' argument to that of the so-called 0-hypothesis put forward by Wijnhoven and colleagues.[39] The 0-hypothesis is that the toxic effects of nanosilver are proportional to the activity of free silver ions released by the nanoparticles. The question, then, for FIFRA regulatory purposes is whether sufficient data and information exist to determine whether nanosilver particles enter the body, whether nanosilver releases silver ions, and to what extent the ions can be absorbed.

The SAP released the minutes of its November 3 to 5, 2009, meeting regarding evaluation of the hazard and exposure associated with nanosilver and other nanometal pesticide products. The panel generally concluded that existing models "are not appropriate" for use with silver nanomaterials and "will not accurately predict nanosilver exposure scenarios." The panel stated that it "strongly believe[s] that in addition to current data requirements under [FIFRA], additional assays which compared nanoscale and bulk materials would be most beneficial in addressing" differences in toxicokinetics and toxicodynamics for nanoscale materials.[40] The panel agreed that pesticide products should be tested on a "case-by-case basis," the EPA should use a meta-analysis on the products to better understand trends in life-cycle analyses, and "close attention" should be given to products that claim a nonionic mode of action as an antimicrobial agent. The panel outlined detailed and extensive research needs that the EPA should consider requiring to develop information on the potential risks of nanoscale silver. Other countries are considering similar issues.[41]

STATE AND LOCAL REGULATORY INITIATIVES

State and local governments are also beginning to focus on the regulation of nanoscale materials from a governance perspective. On December 12, 2006, the Berkeley, California, City Council unanimously approved a proposal to require businesses to report nanoparticles being used, provide available toxicologic information, and outline measures for safe handling of the materials. All businesses that manufacture or use nanoparticles must submit a written report of the current toxicology of the nanomaterials reported and methods for safely handling, monitoring, containing, disposing, and tracking the inventory.[42]

On January 8, 2007, the City Council of Cambridge, Massachusetts, asked the Cambridge Public Health Department to review the Berkeley ordinance and recommend a similar statute for Cambridge. The Cambridge Chamber of Commerce solicited the view of companies, laboratories, and other organizations active in the manufacture, research, and/or use of nanomaterials to ensure full industry participation in the City's review of the need for regulation and the possible development of statutes to that end. On July 28, 2008, the City Council voted to accept a set of recommendations for a municipal health and safety policy on nanomaterials. The recommendations were set forth in a report prepared by the Cambridge Public Health Department and the Cambridge Nanomaterials Advisory Committee.[43][§] Cambridge thus became the second city in the United States to take municipal action on nanomaterials.

More recently, the California Department of Toxic Substances Control (CDTSC) has been obtaining data and information on nanoscale materials. In a January 22, 2009, letter to stakeholders, the CDTSC announced that it is requiring the submission of data "regarding analytical test methods, fate and transport in the environment, and other relevant information from manufacturers of carbon nano-tubes."[44] The CDTSC stated that the term "manufacturers" includes persons and businesses that produce CNTs in California or import CNTs into California for sale and identified manufacturers that produce or import CNTs in California, including academic institutions doing CNT research and those manufacturers that are involved in producing or importing CNTs in their chemical form. The CDTSC stated that the call-in "also includes companies outside California who may export carbon nanotubes into the State. Initially, we have not included manufacturers who import products containing carbon nanotubes; however, we may expand the list of manufacturers in the future to include product manufacturers."[45]

On January 25, 2010, the CDTSC posted on its website responses from 17 companies that received a formal information request letter regarding CNTs.[46] The CDTSC also listed the companies that failed to respond by the January 22, 2010, due date and issued follow-up letters to nine companies that reportedly failed to respond timely on February 16, 2010. The CDTSC also expressed interest in other nanomaterials. In 2010, the CDTSC stated that it will be focusing on other nanoscale materials, including nanometal oxides such as nanoscale titanium dioxide and nanoscale zinc oxide, and nanometals such as nanosilver and nano zerovalent iron.

Risk Management and Product Stewardship Strategies

Given the relative paucity of legal and regulatory standards specific to nanomanufacturing operations, nanotechnology businesses and other stakeholders have devoted considerable effort to developing alternatives to traditional "command and control" regulation to identify and manage risk. These strategies are discussed below.

Key Industry Standard-Setting Initiatives

Several major efforts are under way to develop standards involving nanotechnology. International Organization for Standardization Technical Committee 229 is preparing international consensus standards on several aspects of nanotechnology, including vocabulary, terms, and definitions; measurement and metrology; and EHS.[47]

ASTM International is working on a similar set of standards.[48] ASTM International Committee E 56 on Nanotechnology is developing standards and guidelines for nanotechnology, specifically including terminology and nomenclature; characterization, environmental, and occupational safety and health; international law and intellectual property;

[§] The recommendations call for Cambridge to take the following steps: establish an inventory of facilities that manufacture, process, handle, or store engineered nanoscale materials; offer technical assistance to help firms and institutions evaluate their existing health and safety plans; offer up-to-date health information to residents on products containing nanomaterials; track rapidly changing developments in research; track the evolving status of regulations and best practices concerning engineered nanoscale materials; and report back to the City Council every other year on the changing regulatory and safety landscape.

liaison and international cooperation; and standards of care and product stewardship.

Key Government-Led Initiatives

In addition to the NMSP discussed above and related voluntary initiatives mushrooming internationally, the OECD is engaged in a robust nanotechnology program. The OECD is an intergovernmental organization that includes 30 member countries, of which the United States is one, and maintains relationships with 70 others.[49]

Two OECD group activities are developed relevant to nanoscale materials. As noted above, in September 2006, the OECD established the WPMN. The chemicals sector is the principal focus of the WPMN, and the EPA, FDA, OSHA, and NIOSH are engaged in the WPMN. Current projects involve nine "sector groups" working on separate projects, including generating an EHS database, developing an EHS research strategy, evaluating existing testing guidelines, safety testing a representative set of manufactured nanomaterials, cooperating and exchanging information on voluntary reporting and regulatory schemes, performing risk assessment and exposure measurement, and developing a life-cycle assessment of nanoscale materials.

In March 2007, the OECD created the Committee on Scientific and Technological Policy, which focuses on considering applications of nanotechnologies. The committee's primary objective is to promote international cooperation that facilitates research, development, and the responsible commercialization of nanotechnology in member countries.

Key Private-Sector Stewardship Initiatives

In June 2007, the Environmental Defense Fund (EDF) and DuPont formally announced the release of their joint effort, the *Nano Risk Framework*. The framework is rapidly becoming the standard for measuring best management practice in the nanotechnology industry. The framework defines "a systematic and disciplined process for identifying, managing, and reducing potential environmental, health, and safety risks of engineered nanomaterials across all stages of a product's 'lifecycle'—its full life from initial sourcing through manufacture, use, disposal or recycling, and ultimate fate."

The EDF and DuPont began their collaborative effort to develop the framework in September 2005. They released a draft version to the public on February 26, 2007, and received comments from a diverse array of stakeholders—government, academia, public interest groups, and both large and small companies. In addition to considering the various comments, the EDF and DuPont conducted pilot-testing on surface-treated, high-rutile phase titanium dioxide (TiO_2), single- and multi-walled CNTs, and nanosized zerovalent iron (nano-FeO) "to ensure that [the framework] is flexible, practical, affordable, and effective."[50] The final document issued today "offers guidance on the key questions an organization should consider in developing applications of nanomaterials, and on the information needed to make sound risk evaluations and risk-management decisions." The framework is intended to support ongoing regulatory initiatives, not replace them.

The EDF and DuPont believe that the framework, which is aimed primarily at organizations, both private and public, that are actively working with nanomaterials and developing associated products and applications, will help users organize and evaluate currently available information; assess, prioritize, and address data needs; and communicate clearly how risks are being mitigated. Ultimately, the EDF and DuPont "believe that the adoption of the *Framework* can promote responsible development of nanotechnology products, facilitate public acceptance, and support the formulation of a practical model for reasonable government policy on nanotechnology safety."

The framework consists of six distinct steps and is intended to be used iteratively as stages of development advance and new information becomes available. The six steps are as follows:

Step 1: describe material and application. The first step is to develop a general description of the nanomaterial and its intended uses, based on information in the possession of the developer or in the literature. The user also identifies analogous materials and applications that may help fill data gaps in this and other steps.

Step 2: profile life cycle(s). Step 2 defines a process to develop three sets of profiles: the nanomaterial's properties, its inherent hazards, and associated exposures throughout the life cycle. The user considers the nanomaterial's full life cycle, from material sourcing, through production and use, to end-of-life disposal or recycling. The user considers how the material's properties, hazards, and exposures may change during that life cycle.

Step 3: evaluate risks. In this step, all of the information generated in the profiles is reviewed to identify and characterize the nature, magnitude, and probability of risks presented by the nanomaterial and its anticipated application. The user considers gaps in the life-cycle profiles, prioritizes those gaps, and determines how to address them, either by generating data or by using, in place of such data, "reasonable worst case" assumptions or values.

Step 4: assess risk management. In the fourth step, the user evaluates the available options for managing the risks identified in step 3 and recommends a course of action. Options include engineering controls, personal protective equipment, risk communication, and product or process modifications.

Step 5: decide, document, and act. In step 5, the user consults with the appropriate review team and decides whether or in what capacity to continue development and production. Consistent with transparent decision making, the user documents those decisions and their rationale and shares appropriate information with the relevant internal and external stakeholders. A worksheet is provided in the appendix for documenting information, assumptions, and decisions.[51]

Step 6: review and adapt. Through regularly scheduled and triggered reviews, the user updates and reexecutes the risk evaluation, ensures that risk management systems are working as expected, and adapts those systems in the face of new information or new conditions. Reviews may be prompted by development milestones, changes in production or use, or new hazard or exposure data. As in step 5, the user not only documents changes, decisions, and actions but also shares appropriate information with relevant stakeholders.

A second initiative that is credible and well thought out is the *Responsible NanoCode*.[52] Britain's Royal Society, the Nanotechnology Industries Association, Insight Investment, and the UK government-sponsored Nano-technology Knowledge Transfer Network collaborated on the proposed code. The objective of this "principles-based" voluntary code of conduct is to encourage industries, retailers, universities, research institutes, and other public or privately funded bodies involved in developing, manufacturing, and selling products of nanotechnology to adhere to seven principles to demonstrate responsible governance. The seven principles are as follows:

- Principle One—Each organization should ensure that responsibility for guiding and managing its involvement with nanotechnologies resides with the Board or governing body.
- Principle Two—Each organization should proactively engage with its stakeholders and be responsive to their views in its development or use of products using nanotechnologies.
- Principle Three—Each organization should identify and minimize sources of risk for workers handling products using nanotechnologies, at all stages in the production process or in industrial use, to ensure high standards of occupational health and safety.
- Principle Four—Each organization should carry out thorough risk assessments and minimize any potential public health, safety, and environmental risks relating to its products using nanotechnologies.
- Principle Five—Each organization should consider and respond to any social and ethical implications and impacts during the development or sale of products using nanotechnologies.
- Principle Six—Each organization should adopt responsible practice in the sales and marketing of products using nanotechnologies.
- Principle Seven—Each organization should engage with suppliers and/or business partners to encourage and stimulate their adoption of the Code and so ensure its own ability to fulfill its Code commitments.

Code proponents launched a consultation period in the United States on October 9, 2007, which ended on November 12, 2007. On May 13, 2008, the Working Group of the Responsible Nano Code signed off on the code.[53]

A third private sector initiative is the Good-NanoGuide, a collaboration platform designed to enhance the ability of experts to exchange ideas on how best to manage nanomaterials in occupational settings.[54] The Good-NanoGuide's beta sponsors include the NIOSH, International Council on Nanotechnology (ICON), Nanotech BC, Nano Alberta, Health Canada, Nano Quebec, and IRSST (a scientific research organization in Quebec). The GoodNanoGuide has three goals: develop and launch a protected Internet site on occupational practices for the safe handling of nanomaterials using a Wiki software platform; create a process wherein multiple stakeholders within the international community contribute, share, and discuss information related to occupational safety; and establish a modern, interactive forum that fills the need for up-to-date information and remains current as new practices develop.

Although the GoodNanoGuide reflects the global dialogue under way regarding the effect that nanotechnologies may have on human health, the environment, and society in general, the GoodNanoGuide is not meant to address or resolve such issues. Instead, according to the ICON, "it assumes that someone, somewhere in the world is likely to be working on nanomaterials as this debate continues, and thus endeavors to provide information to that person about current good practices to make sure appropriate safeguards are in place as that person works on nanomaterials in an occupational setting."[55]

Finally, the GoodNanoGuide is open for everyone to review. To ensure the dependability of the good practices reported, the GoodNanoGuide is a protected site in which contributions are limited to those individuals who have become GoodNanoGuide members.

CONCLUSION

As with the other commercial, legal, and governance issues, businesses engaged in the production and use of nanoscale materials must look to existing laws and regulations, voluntary and stewardship initiatives, and best industry practices to avoid liability. At the same time, these entities must discover new ways to proceed in an arena where the state of knowledge is still catching up to entrepreneurial initiatives.

REFERENCES

1. National Nanotechnology Initiative. What is nanotechnology? Available at: http://www.nano.gov/html/facts/whatIsNano.html (accessed May 31, 2010).
2. Lux Research. The Nanotech report. 4th ed. 2006. Available at: http://www.luxresearchinc.com/pdf/TNR4_TOC.pdf (accessed July 2, 2010).
3. Oberdorster G, Oberdorster E, Oberdorster J. Nanotechnology: an emerging discipline evolving from studies of ultrafine particles. Environ Health Perspect 2005;113:823.
4. Environmental Protection Agency. Control of Nanoscale Materials under the Toxic Substances Control Act. EPA Web site on Nanotechnology. Available at: http://www.epa.gov/oppt/nano/ (accessed July 2, 2010).
5. Food and Drug Administration. Nanotechnology. Available at: http://www.fda.gov/nanotechnology (accessed July 2, 2010).
6. Occupational Safety and Health Administration. Nanotechnology. Available at: http://www.osha.gov/dsg/nanotechnology/ nanotechnology.html (accessed July 2, 2010).
7. Centers for Disease Control and Prevention. Nanotechnology at NIOSH. Available at: http://www.cdc.gov/niosh/topics/nanotech/ (accessed July 2, 2010).
8. National Toxicology Program, National Institute of Environmental Health Sciences, National Institutes of Health. NTP Nanotechnology Safety Initiative. Available at: http://ntp.niehs.nih.gov/?objectid=7E6B19D0-BDB5-82F8-FAE73011304F542A (accessed July 2, 2010).
9. US Department of Defense. Department of Defense Web site on Nanotechnology. Available at: http://www.defenselink.mil (accessed July 2, 2010).
10. Environmental Protection Agency. STAR grants. Available at: http://www.epa.gov/ncer/grants/ (accessed July 2, 2010).
11. Project on Emerging Nanotechnologies, Woodrow Wilson Center for International Scholars.

A nanotechnology consumer products inventory. Available at: http://www.nanotechproject.org/inventories/ consumer/ (accessed May 31, 2010).
12. Nano particles were not the cause of health problem triggered by sealing sprays! Available at: http://www.bfr.bund.de/cd/7842 (accessed May 31, 2010).
13. Roberts K. Nanotechnology and the future of litigation. ABA Section of Litigation Magazine 2010; 31(2): 6–8.
14. Letter from shareholder rights activists to then President-Elect Obama. Available at: http://icbn.org/Obamaletter.pdf (accessed July 22, 2010).
15. Proxy Review 2009. Available at: http://www.asyousow.org/publications/ProxyReview_2009.pdf (accessed May 31, 2010).
16. Rizzuto R. Increase expected in shareholder resolutions urging disclosure of nanomaterials, policies. Bureau of National Affairs. Daily Environmental Report 2009; Jan 15: A-6.
17. Investor Environmental Health Network. Resolutions introduction. Available at: http://www.iehn.org/resolutions.introduction.php (accessed May 13, 2010).
18. 74 Federal Register 68215 (December 23, 2009).
19. 75 Federal Register 3462 (January 21, 2010).
20. 75 Federal Register 29754 (May 27, 2010).
21. Swiss Re. Nanotechnology small matter, many unknowns. Available at: http://media.swissre.com/documents/nanotechnology_ small_matter_many_unknowns_en.pdf (accessed May 31, 2010).
22. Lexington Insurance. Available at: http://www.lexingtoninsurance.com/documents/lexHSLexNanoShield.pdf (accessed May 31, 2010).
23. Nanoscale Materials Stewardship Program. Concept Paper: 1–3.
24. Environmental Protection Agency. Nanoscale Materials Stewardship Program interim report. January 2009. Available at: http://www.epa.gov/oppt/nano/nmsp-interim-report-final.pdf (accessed May 31, 2010).
25. 75 Federal Register 5546 (February 3, 2010).
26. Qassim A. Study says U.S., Europe need to strengthen reporting requirements for nanomaterials. Bureau of National Affairs. Daily Environmental Report 2009; Sept 11: A-9.
27. Notice of availability. 73 Federal Register 4861 (January 28, 2008).
28. 73 Federal Register 64946 (October 31, 2008).
29. 73 Federal Register 25696 (May 7, 2008).
30. 73 Federal Register 65743 (November 5, 2008).
31. 74 Federal Register 280 (January 5, 2009).
32. 74 Federal Register 29982 (June 24, 2009).
33. 74 Federal Register 42177 (August 21, 2009).
34. 75 Federal Register 5546 (May 31, 2010).
35. Jordan W. PowerPoint presentation to the Pesticide Program Dialogue Committee. April 29, 2010. Available at: http://www.epa.gov/pesticides/ppdc/2010/april2010/session1-nanotec.pdf (accessed May 31, 2010).
36. 73 Federal Register 69644 (November 19, 2008).
37. Environmental Protection Agency. Pesticide issues in the works: nanotechnology, the science of small. July 22, 2008. Available at: http://www.epa.gov/pesticides/about/intheworks/nanotechnology.htm (accessed May 31, 2010).
38. Environmental Protection Agency. FIFRA Scientific Advisory Panel Background Paper: Evaluation of Hazard and Exposure Associated with Nanosilver and Other Nanometals. November 3–6, 2009. Available at: http://www.epa.gov/scipoly/sap/meetings/2009/ novemsere/110309agenda.pdf (accessed August 17, 2010).
39. Wijnhoven S, Peijnenburg W, Herberts C, et al. Nanosilver–a review of available data and knowledge gaps in human and environmental risk assessment. Nanotoxicology 2009;3:109–138.
40. Environmental Protection Agency. SAP Minutes No. 2010-01. A Set of Scientific Issues Being Considered by the Environmental Protection Agency Regarding: Evaluation of the Hazard and Exposure Associated with Nanosilver and Other Nanometal Pesticide Products. November 3–5, 2009 FIFRA Scientific Advisory Panel Meeting. Available at: http://www.epa.gov/scipoly/sap/meetings/2009/november/110309anseetingminutes.pdf
41. UK Advisory Committee on Hazardous Substances report on nanosilver. Available at: http://www.defra.gov.uk/environment/quality/chemicals/achs/documents/achs-report-nanosilver.pdf (accessed May 31, 2010).
42. Berkeley Municipal Code. Sections 15.12.040,050. Available at: http://www.ci.berkeley.ca.us/citycouncil/ordinances/2006/6960.pdf (accessed May 31, 2010).
43. Cambridge Nanomaterials Advisory Committee. Cambridge Public Health Department. Recommendations for a Municipal Health & Safety Policy for Nanomaterials: a Report to the Cambridge City Manager. July 2008. Available at: http://www.cambridgepublichealth.org/publications/July_17_08_Nano_ Recommendations.pdf (accessed May 31, 2010).

44. California Department of Toxic Substances Control. Available at: http://www.dtsc.ca.gov/TechnologyDevelopment/Nanotechnology/upload/Formal_AB289_Call_In_Letter_CNTs.pdf (accessed May 31, 2010).
45. California Department of Toxic Substances Control. Available at: http://www.dtsc.ca.gov/TechnologyDevelopment/Nanotechnology/upload/AB289_CNT_Contact_List.pdf (accessed May 31, 2010).
46. Chemical Information Call-in, DTSC. Available at: http://www.dtsc.ca.gov/PollutionPrevention/Chemical_Call_In.cfm (accessed May 31, 2010).
47. International Organization for Standardization. Available at: http://www.iso.org/iso/iso_catalogue.htm (accessed May 31, 2010).
48. ASTM International. Available at: http://www.astm.org (accessed May 31, 2010).
49. Organization for Economic Cooperation and Development. Available at: http://www.oecd.org (accessed May 31, 2010).
50. Environmental Defense Fund and DuPont. Nano Risk Framework. Available at: http://nanoriskframework.com/page.cfm?tagID=1095 (accessed May 31, 2010).
51. Environmental Defense Fund and DuPont. Nano Risk Framework Case studies. Available at: http://nanoriskframework.org/page.cfm?tagID=1326 (accessed May 31, 2010).
52. Responsible NanoCode: consultation draft–17 September 2007. Available at: http://www.docstoc.com/docs/26560846/ nanotechnology-governance (accessed May 31, 2010).
53. The Responsible NanoCode–update May 2008. Available at: http://responsiblenanocode.org/documents/The Responsible Nano Code Update Announcement.pdf (accessed May 31, 2010).
54. The GoodNanoGuide. Available at: http://www.goodnanoguide.org (accessed May 31, 2010).
55. ICON GoodNanoGuide Statement of Objectives. Available at: http://icon.rice.edu/projects.cfm?doc_id_= 12207 (accessed May 31, 2010).

2.3.3 NANOTECHNOLOGY—A NEW FIELD OF ETHICAL INQUIRY?

ARMIN GRUNWALD

Many people believe that nanotechnology (NT) has the potential to dramatically impact our lives and the future of technology. In this essay, Armin Grunwald, a German physicist and professor of philosophy, discusses a range of ethical issues and threats posed by recent advances in this field.

Although many important technological discoveries have produced both beneficial and negative results, some of them unanticipated, NT has unique properties, which pose radically new challenges. The fact that the building blocks of life operate at the nanoscale raises the question of whether advances in NT will allow for human control of basic life processes. According to Grunwald, NT may require a new "ethics of responsibility" (see Jonas, 1.3.1) and raise new problems for equity and distributive justice.

Advances in NT are likely to increasingly blur the line between technology and organic life. For those with debilitating and life-threatening medical conditions, such techniques could lead to revolutionary new cures. But, as Grunwald points out, we must also address potentially negative effects of such treatments and how we might distribute them. In addition to potentially helping people with serious medical

Source: From "*Science and Engineering Ethics*," Volume 11, Issue 2, 2005, Grunwald, "Nanotechnology—A New Field of Ethical Inquiry?" pp. 187–201. Copyright © 2005 by Opragen Publications. Reprinted with kind permission from Springer Science+Business Media B.V.

conditions, NT might also be used to enhance human capabilities. This possibility raises serious ethical concerns, both about attempting to improve and control human life, but also about who might and might not receive—or be able to afford—these enhancement technologies.

Because of the microscopic size of NT devices, questions of privacy and security need also to be considered. It is likely that it will be possible to gather information passively using NT in ways it is difficult to imagine today. Thus NT, along with other new technologies, has the potential to greatly extend the surveillance capabilities of the state, corporations, and others. How we handle such ethical challenges may be vital to the shape of the human future.

FOCUS QUESTIONS

1. In what major ways is NT different than previous technologies, and why is it seen by some as posing unique threats and raising unprecedented ethical dilemmas?
2. How do the concerns raised here about enhancing humans compare with those discussed by Sandel (2.5.3) and Fukuyama (2.5.1)?
3. What issues does Grunwald identify with regard to the "distributive justice" and "sustainability" of NT advances? Discuss how these concerns are either unique to this technology or apply to others as well.

KEYWORDS

converging technologies, cyborgs, enabling technology, enhancement of human beings, equity, ethics of responsibility, humanoid "lab-on-a-chip," liberal eugenics, microimplants, nano-ethics, nanoparticles, nanoscale, nanotechnology, precautionary principle, privacy, RFID-technology, risks, sustainability side effects, technological determinism

1. INTRODUCTION

In view of the revolutionary potential attributed to nanotechnology,[1,2] it is not surprising that this technology has found great interest in the media and in the public at large. Ethical, legal, and social implications (ELSI) are already being elaborated by commissions and study groups. Ethical reflection on nanotechnology has, in fact, already coined new terms such as "nano-ethics,"[3,4] but has, to date, hardly accomplished more than to proclaim a need for ethics in and for nanotechnology.[5] The ethical aspects discussed in the—to the present, few—treatises are much rather evidence of a tentative approach to a relatively new field of science and technology than of systematic analysis. Criteria for determining why certain topics, such as self-replicating nanorobots or nanoparticles, should be ethically relevant are not given. The normative fuzziness of unclear criteria joins the cognitive fuzziness caused by a lack of knowledge of the real possibility of the technical innovations concerned. The result, as the working hypothesis of this contribution has it, discloses, above all, a diffuse uneasiness regarding rapid scientific advances in nanotechnology.

In view of this situation, the purpose of this contribution consists primarily in studying current and foreseeable developments in nanotechnology from the viewpoint of ethics. Which developments are ethically relevant? Are there ethical questions which have already been answered by current discussions in the ethics of technology or in bioethics? Are there developments which pose completely new questions? To this end, it is necessary—besides inquiry into the present state of the discussion—to look for criteria for deciding when something is ethically relevant (sec. 2). These criteria are then applied to the field of nanotechnology, and the ethical challenges are "mapped out" (sec. 3). This is done on the basis of an evaluation of the current literature, and of

pertinent references on the Internet. After answering the question on the necessity of nano-ethics (sec. 4), comments are subsequently made on the role to be played by ethics in the further development of nanotechnology.

2. ETHICS FOR NANOTECHNOLOGY—RELEVANCE CRITERIA

It has been controversial for quite some time whether science and engineering have any morally relevant content at all (and could therefore be a subject of ethical reflection). Into the 1990's, technology was held to be *value-neutral*. In numerous case studies, however, the normative background of decisions on technology (even those made in the laboratory) have since been recognized and made a subject of reflection.[6] The basic assumption of this development is not to see technology solely as a sum of abstract objects or processes, but to take its embeddedness in societal processes seriously. Technology is not nature, and does not originate of itself, but is consciously produced to certain ends and purposes—namely, to bring something about which wouldn't happen of itself. Technology is therefore always embedded in societal goals, problem diagnoses, and action strategies. In this sense, there is no "pure" technology, *i.e.* a technology completely independent of this societal dimension.

Normative aspects of science and technology, in a morally pluralistic society, unavoidably lead to societal debates at the least, and often also to conflicts over technology. As a rule, what is held to be desirable, tolerable, or acceptable is controversial in society. Open questions and conflicts of this type in the context of science and technology are the point of departure for the ethics of technology.[7] Technology conflicts are, as a rule, not only conflicts over technological means (e.g., in questions of efficiency), but are also conflicts over visions of the future, of concepts of humanity, and on views of society. The role of the ethics of technology consists of analysis of the normative structure of technology conflicts and the search for rational, argumentative, and discursive methods of resolving them.

Even if technology is basically beset with values, most of the technically relevant decisions can be classified as the "standard case" in the following sense:[7,8] they don't subject the normative aspects of the basis for the decision (criteria, rules or regulations, goals) to specific reflection, but assume them to be given for the respective situation, and accept the frame of reference they create. In such cases, no explicit ethical reflection is, as a rule, necessary, even if normative elements self-evidently play a vital role in these decisions—the normative decision criteria are clear, acknowledged, and unequivocal. It is then out of the question that this could be a case of conflict with moral convictions or a situation of normative ambiguity—the information on the normative framework can be integrated into the decision by those affected and by those deciding on the basis of axiological information, without analyzing it or deliberating on it. The (national and international) legal regulations, the rules of the relevant institutions (e. g., corporate guidelines), where applicable, the code of ethics of the professional group concerned, as well as general societal usage, are elements of this normative framework.

The requirements to be set in the normative framework in order that it can be accepted as a "standard case" can be operationalized according to the following criteria:[7,8]

- *Pragmatic completeness*: the normative framework has to treat the decision to be made fully with regard to normative aspects;
- *Local consistency*: there must be a "sufficient" measure of consistency between the normative framework's elements;
- *Unambiguity*: among the relevant actors, there must be a sufficiently consensual interpretation of the normative framework;
- *Acceptance*: the normative framework must be accepted by those affected as the basis for the decision;
- *Compliance*: the normative framework also has to be complied with in the field concerned.

There are such well-established "standard" situations in many fields (e.g., in decision-making processes in public administration or in private businesses). Technical innovations and scientific

progress, however, can challenge such situations by presenting new questions or by shaking views previously held to be valid. This is then the entry point for ethical reflection in questions of science and engineering, for the explicit confirmation, modification, or augmentation of the normative framework.

Whether there are new challenges for ethics in nanotechnology and what they might be, will be investigated against this background below. In order to clarify the question of ethical relevance, an examination of the topics treated in previous papers judged on the basis of systematic criteria for the ethical relevance of science and technology was carried out.

The guiding questions in this study were:

- Which are the ethical aspects in the sense defined above, especially which are the genuine ethical aspects in the subjects for ethics in nanotechnology named in current publications?
- Is relevant and sufficiently evident knowledge available for the assessment of scientific and technical developments in nanotechnology or of their use (as addressed with regard to ethics)?
- Which of the ethical aspects of nanotechnology are specific for nanotechnology?

This method permits a clear determination of the questions to which ethical reflection can make a real contribution. It then also permits an estimation about to what extent an independent "nanoethics" would be justified. The resulting "map" of ethical aspects described below is a survey of the use of nanotechnology as regards aspects of content, and includes partially overlapping fields of content. These fields have the fact in common that nanotechnology will play a constitutive role in the foreseeable advances in their areas.

3. ETHICAL ASPECTS OF NANOTECHNOLOGY

In practical philosophy and in ethics, nanotechnology has to date seldom been made a subject of discussion. [NT = nanotechnology]: "While the number of publications on NT per se has increased dramatically in recent years there is very little concomitant increase in publications on the ethical and social implications to be found."[8] Some postulate a need for ethics and make reference, above all, to remote visions (for example, the abolition of aging or "runaway nanorobots"[8,9]). Certain terms, such as privacy, man-machine relationship, or equity are often cited (e.g.,[10,11]). Systematic studies, which could do justice to the diversity and breadth of nanotechnology, haven't yet been presented. At the same time, concern over the detrimental consequences of this deficit grows: "The lack of dialogue between research institutes, granting bodies and the public on the implications and directions of NT may have devastating consequences, including public fear and rejection of NT without adequate study of its ethical and social implications."[4]

3.1 Nanoparticles–Opportunities versus Risks

A vast potential market for nano-based products is seen in the field of new materials.[1] By means of an admixture or specific application of nanoparticles, some new properties of materials can be brought about, for instance, in surface treatment. If making new technical functions and performance factors is a common motivation for scientific and technological progress, ethical questions on these advances pose themselves as the "classical" question of the possible *side effects* of these developments.[1]

Artificially produced nanoparticles can be disseminated into the environment or enter the human body as emissions during production or by the daily use of products including nanoparticles. Nanoparticles could eventually be transported as aerosols over great distances and be distributed diffusely. They could enter the human body by way of the lungs, through the skin, or the digestive tract. How they interact in spreading, their impact on health and on the environment, in particular, potential long-term effects, are at present almost unknown. This applies also and above all for substances which don't occur in the natural environment, such as fullerenes or nanotubes. As far as the potential proliferation of nanoparticles is concerned, aspects such as

mobility, reactivity, persistence, lung infiltration, solubility in water, etc., have to be taken into consideration.

Questions of toxicity for the environment and for humans, on nano-material flow, on their behavior in spreading throughout various environmental media, on their rate of degradation, and their consequences for the various conceivable targets are, however, not ethical questions. In these cases, the pertinent empirical-scientific disciplines, such as toxicology or environmental chemistry, are competent. Their results would be of interest from the viewpoint of ethics as soon as the available empirical facts can be consulted to establish which practical consequences they have for working with nanoparticles:

- What follows from our present lack of knowledge about the possible side effects of nanoparticles? One radical consequence, namely, a moratorium on putting nanoparticles onto the market, as it would probably follow from Hans Jonas's ethics of responsibility[12] according to the priority of the negative prognosis, has already been demanded.[13]
- More generally formulated: is the precautionary principle relevant in view of a lack of knowledge,[14,15] and what would follow from the answer?
- Which role do the—doubtlessly considerable—opportunities of nanoparticle-based products play in considerations of this sort? According to which criteria can judgment of benefits and hazards be based when the benefits are (relatively) concrete, but the hazards are hypothetical?
- Are comparisons of the risks of nanoparticles with types of risk known from other areas possible, in order to learn from them? Can criteria for assessing the nanoparticle-risks be gained from experience in developing new chemicals or medicines? Which normative premises enter into such comparisons?
- Is the discussion on threshold values and environmental standards[16] transferable to nanoparticles? What about the acceptability or tolerability of risks? Which methods of determining critical values come into question, and how do they relate to ethical questions or questions of democratic theory?

The contribution of ethics to this subject therefore lies in a value judgment of the situation (relationship of reliable knowledge to the degree of uncertainty), to clarification of the comparability to other types of risk, and by disclosing the normative presuppositions and implications entering into it, as well as to the normative basis for practical consequences. These questions of the acceptability and comparability of risks, the advisability of weighing up risks against opportunities, and the rationality of action under uncertainty are, without doubt, of great importance in nanotechnology. A new field of application is developing here for the ethics of science and technology; the type of questions posed, however, is well-known from other discussions of risks (exposure to radiation; chemicals).

3.2 Sustainability: Just Distribution of Opportunities and Risks

Possible side effects of nanotechnology of a completely different type result from considerations of equity. In particular, in connection with sustainability[17] are ethical questions of the *distribution* of the opportunities for the use of as well as the spatial and temporal distribution of the opportunities and risks of nanotechnology:[18] "Nanotech offers potential benefits in areas such as biomedicine, clean energy production, safer and cleaner transport and environmental remediation: all areas where it would be of help in developing countries. But it is at present mostly a very high-tech and cost-intensive science, and a lot of the current research is focused on areas of information technology where one can imagine the result being a widening of the gulf between the haves and the have-nots."[8] In this respect, we have to distinguish between intragenerational and intergenerational aspects.

The distribution of the use of natural resources between present and future generations belongs to the *intergenerational* aspects. Appreciable relief for the environment is expected from the use of nanotechnology: savings of material resources, reduction of the incidence of environmentally detrimental by-products, improvement

reduction of energy consumption, and the elimination of ecologically deleterious materials from the environment.[1] Decisive for the assessment of nanotechnology or of product lines based on nanotechnology from the viewpoint of sustainability is that technology accumulates positive and negative contributions to sustainability throughout its entire "lifetime," which extends from the primary sources of raw materials via transportation and processing to consumption, and finally ends with their disposal as waste.[19] The entire life cycle of a technology is therefore decisive for the assessment of the technology's sustainability. But nanotechnology is at present in many branches still in an early phase of development. For this reason, we can only speak of nanotechnology's *sustainability potentials*.[18] There is no guarantee that technical sustainability *potentials* will turn out to be real contributions to sustainable development. The discussion of sustainability potentials can, however, be employed constructively with regard to technological development, if further development is accompanied by ethical reflection on questions of distribution between present and future exploitation of nature.

Intragenerational problems of distributive justice present themselves basically in every field of technical innovation. Because scientific and technical progress requires considerable investments, it usually takes place where the greatest economic and human resources are already available. Technical progress increases existing inequalities of distribution. This can be illustrated on the example of nanotechnology in medicine.[18] Nanotechnology-based medicine will, in all probability, be expensive medicine. Questions of equity and of access to (possible) medical treatments could become urgent in at least two respects: *within* industrialized societies, existing inequalities in access to medical care could be exacerbated by a highly technicized medicine making use of nanotechnology, and—with regard to *less developed societies*—because likewise, already existing and particularly dramatic inequalities between technicized and developing nations could be still further increased. Apprehensions with regard to both of these types of a potential "nano-divide" (after the well-known "digital divide") are based on the assumption that nanotechnology can not only lead to new and greater options for individual self-determination (e. g., in the field of medicine), as well as to considerable improvement of the competitiveness of national economies. Current discussions on distributive justice on the national and on the international level (in the context of sustainability as well) are therefore likely to gain new relevance with regard to nanotechnology.

Both of the aspects described are, however, not really new ethical aspects of technology, but are rather intensifications of problems of distribution already rife. Problems of equity belong in principle to the ethical aspects of modern technology.

3.3 The Private Sphere and Control

Another field regularly mentioned among the ethical aspects of nanotechnology is the threat to privacy through new monitoring and control technologies.[8,20] Nanotechnology offers a range of possibilities for gathering, storing, and distributing personal data to an increasing extent. In the course of miniaturization, a development of sensor and memory technology is conceivable which, unnoticed by its "victim," drastically increases the possibilities for acquiring data.[3] Furthermore, miniaturization and networking of observation systems could considerably impede present control methods and data protection regulations, or even render them obsolete. For the military, new possibilities for espionage are opened.[21] Passive observation of people could, in the distant future, be complemented by actively manipulating them—for instance, if it would be possible to gain direct technical access to their nervous system or brain.

These scenarios are regarded by some to be not only realistic, but even certain: "But what is not speculation is that with the advent of nanotechnology invasions of privacy and unjustified control of others will increase."[3] Underlying this opinion is an only thinly veiled technological determinism: "When new technology provides us with new tools to investigate and control others we will use them."[3]

Within the private sphere, health is a particularly sensitive area. The development of small

analyzers—the "lab on a chip"—can make it possible to compile comprehensive personal diagnoses and prognoses on the basis of personal health data. Stringent standards for data protection and for the protection of privacy therefore have to be set. Without sufficient protection of their private sphere, people are rendered manipulable, their autonomy and freedom of action are called in question.[1] The "lab-on-a-chip" technology can facilitate not only medical diagnoses, but can also make fast and economical comprehensive screening possible (*op. cit.*). The rapid decoding of complete individual genetic dispositions can come within the reach of normal clinical work or of clinic-external services. Everyone could let him- or herself be tested, for example, for genetic dispositions for certain disorders—or could be urged by his/her employer or insurance company to do so. In this manner, individual persons could find themselves put under social pressure, their freedom of action would be impaired. In addition, it would have to be clarified, how one has to deal with results which possibly depress the patients affected over longer periods of time, fearing the impending outbreak of a serious disease (which possibly never occurs). How to deal with—presumably sometimes considerable—uncertainties of diagnoses is, in this connection, also an important aspect.

Questions of privacy, of monitoring and controlling people are doubtlessly ethically relevant. The current discussions on restricting civil liberties for the sake of combating terrorism are taking place against a background in which normative criteria play a pivotal role, and in which there are conflicts over this subject. On the other hand, the *political* dimension of these questions stands prominently in the foreground, for instance, on estimations of the immediacy of the danger and on statements on the assumed problem-solving capacity of measures proposed. If it should come to an endangerment of the private sphere through nanotechnology, the ethical dimension would be more tangible. And nonetheless, it would even in this case be more likely to be kept in the background, while the central discus-

with the context of data protection, in which we have already gained considerable experience.

But all of these questions of monitoring and of data protection are not posed exclusively by nanotechnology. Even without nanotechnology, observation technologies have reached a remarkable stage of development which poses questions on the preservation of the private sphere. Even today, so-called smart tags, based on RFID-technology (Radio Frequency Identification), are being employed for access control, as ticketing, e. g., in public transportation and in logistics. These objects have at present a size of several tenths of a millimetre in each dimension, so that they are practically unnoticeable to the naked eye. Further miniaturization will permit further reduction in size and the addition of more functions—without nanotechnology being needed—but nanotechnology will promote and accelerate these developments.

The ethically relevant questions on a right to know or not to know, on a personal right to certain data, on a right to privacy, as well as the discussions on data protection and on possible undesirable inherent social dynamisms, and, in consequence, of a drastic proliferation of genetic and other tests, have been a central point in the bio- and medical-ethical discussions for quite a while. Nanotechnological innovations can accelerate or facilitate the realization of certain technical possibilities, and therefore increase the urgency of the problematics of the consequences; in this area, however, they don't give rise to qualitatively new ethical questions.

3.4 Crossing the Border between Technology and Life

Basic life processes take place on a nano-scale, because life's essential building-blocks (such as proteins, for instance) have precisely this size. By means of nanotechnology,[2] biological processes are made nanotechnologically controllable. Molecular "factories" (mitochondria) and "transport systems," which play an essential role in cellular metabolism, can be models for controllable bio-nanomachines. Nanotechnology on this level

intermeshing of natural biological processes with technical processes seems to be conceivable. The classical barrier between technology and life is increasingly being breached and crossed. One speaks of elements of living organisms in classical mechanics' language: as factories, rotors, pumps, and reactors.

This is, at first glance, a cognitively and technically extremely interesting process, with a great deal of promise. The technical design of life processes on the cellular level, direct links and new interfaces between organisms and technical systems portend a new and highly dynamic scientific and technological field. Diverse opportunities, above all, in the field of medicine, stimulate research and research funding.[1]

New ethical aspects are certainly to be expected in this field. Their concrete specification, however, will only be possible when research and development can give more precise information on fields of application and products. The corresponding discussions of risks could have structural similarities to the discussion on genetically modified organisms. It could come to discussions about safety standards for the research concerned, about "field trials," and release problems. The danger of misuse will be made a topic of debate, such as, for example, technically modifying viruses in order to produce new biological weapons. A wide range of future ethical discussions is opening, for which at present there is insufficient practical background for concrete reflection. In this field of nanotechnology (not in that of the nanoparticles, [sec. 3.1]), similar developments concerning resistance in society could be feared as in the case of genetic engineering.

A new area that is practically as well as ethically interesting consists of making direct connections between technical systems and the human nervous system. There is intensive work on connecting the world of molecular biology with that of technology. An interesting field of development is nanoelectronic neuro-implants (neurobionics), which compensate for damage to sensory organs, or to the nervous system, or increase the performance capacity of these organs and broaden the spectrum of human perception. Microimplants could restore the functions of hearing and eyesight. Even today, simple cochlear or retina implants, for example, can be realized. With progress in nanoinformatics, these implants could approach the smallness and capabilities of natural systems. Because of these undoubtedly positive goals, ethical reflection could, in this case, concentrate above all on the definition and prevention of misuse. Technical access to the nervous system, because of the possibilities for manipulation and control which it opens, is a particularly sensitive issue.

Extrapolating these lines of development into the realm of speculation, the convergence of technology and humanity, the conceivability (in the sense of a pure thought possibility) of "cyborgs" as technically enhanced humans or humanoid technology could be problematized.[1] Developments of this type raise the question of humanity's self-concept, which is of great ethical relevance. In nanotechnological visions, aspects repeatedly occur which blur the boundary between what human beings are and what they create with the help of technical achievements and applications. Such visions pose the question, to what extent technical or partly technical can partly biologically-constructed man-machine chimeras lay claim to the status of a person. An entire spectrum of anthropological and ethical questions follows from out of this question aspect. In some US-American churches—in contrast to the Transhumanism discussion—precisely this aspect of nanotechnology is a central theme. Nanotechnology acts as a playing-field on which various philosophies of life compete.

In spite of the speculative nature of the subject, ethical reflection doesn't seem to be premature. One needn't see the main concentration of nanotechnology's ethical aspects in this area, but one can certainly find enough indications that scientific and technical progress will intensify the—at present, non-existent—urgency of these questions in the coming years. In particular, advances in brain science and developments in "converging technologies"[2] lead to this expectation—and would justify ethical reflection "in advance." But the insight that the questions raised aren't really specific for nanotechnology also applies here. Since the 1980s, these subjects

have been repeatedly discussed in the debates on artificial intelligence and on artificial life.

3.5 The Improvement of the Human Being

With the transgression of the boundary between technology and living beings from both sides, our understanding of what distinguishes human beings and which relationship mankind assumes to be its natural physical and psychic constitution is also called into question. Within the tradition of technical progress that has, at all times, transformed conditions and developments—which, until then, had been taken as given, as unalterable fate—into influenceable, manipulable, and formable conditions and developments, the human body and its psyche are rapidly moving into the dimension of the Formable. The vision of "enhancing human performance" has been conjured up, above all, in the field of "converging technologies."[1]

If the wish to "improve humanity"—probably in view of the experience of the cultural and social deficits of "real human beings"—has been expressed often in history, the approach of bringing about these improvements by technical means—assuming that this is realizable—is apparently new (science fiction, as a rule, doesn't lay any claim to future realizability). Formerly, improvement utopias were founded rather on the basis of "soft" methods—above all, upbringing and education. Technology had its place outside of human beings, as a means of augmenting mankind's capacities for action. Technical enhancement of human beings themselves—if this would be possible at all—would in any case pose a series of new ethical questions.

Nanotechnology, in combination with biotechnology and medicine, opens perspectives for fundamentally altering and rebuilding the human body. At present, research is being done on tissue and organ substitution, which could be realized with the help of the nano- and stem cell technologies. Nanoimplants would be able to restore human sensory functions or to complement them, but they would also be able to influence the central nervous system. While the examples of medical applications of nanotechnology cited remain within a certain traditional framework—because the purpose consists of "healing" and "repairing" deviations from an ideal condition of health, which is a classical medical goal—chances (or risks) of a remodeling and "improvement" of the human body are opened up. This could mean extending human physical capabilities, e.g., to new sensory functions (for example, broadening the electromagnetic spectrum the eye is able to perceive). It could, however, also—by means of the direct connection of mechanical systems with the human brain—give rise to completely new interfaces between man and machine, with completely unforeseeable consequences. Even completely technical organs and parts of the body (or even entire bodies) are being discussed, which, in comparison with biological organisms, are supposed to have advantages such as—perhaps—increased stability against external influences.[1]

There are initial anthropological questions of our concept of humanity and of the relationship between humanity and technology.[11] With them, however, and at the same time, the question poses itself, how far human beings *can*, *should*, or *want to* go in remodeling the human body, and to what end(s) this should or could be done. The difference between healing and enhancing interventions is on conceptual grounds—in particular the terms "health" and "illness" haven't been clarified to date—and for practical reasons, a gradual approach would be prudent.[22]

The improvement of mankind could also include the postponement of death. According to the definition of health formulated by the World Health Organization (WHO), according to which health is a "state of complete physical, mental, and social well-being, and not merely the absence of disease or infirmity" (Charter of the WHO), aging could also be interpreted as a disorder. Overcoming aging with the help of nanotechnology would, then, in the sense of medical ethics, be nothing other than fighting epidemics or other diseases. Aging "as a disease" would be combated medically as if it were the flu. Seen against the background of continuing and controversial discussions on the concept of illness in medical theory and in medical ethics, this is,

however, not undisputed. Whether aging as a process and death are, in principle, acknowledged as a predetermined initial condition of human existence, and should only be made subject to medical treatment in their extreme expression, or whether aging and death are seen as conditions which are, whenever possible, to be abolished, depends on fundamental normative presuppositions, which, in view of the conflicts connected with them, are of great ethical relevance.[3]

Possible answers from ethics could, according to the respective school of thought, turn out quite differently. Liberal eugenics based on utilitarian ethics could draw the conclusion that it doesn't acknowledge any difference between therapeutic and enhancing interventions and leave "the choice of purposes of interventions which alter characteristic traits to individual preferences."[22] Kantian ethics would thematize the instrumentalization of human beings, religious morals would bring traditional human self-concepts (which are deeply rooted in culture) as a limited being (in time as well) to bear.

The practical relevance of such ethical questions in view of a—at least in the eyes of some protagonists—possible technical improvement of human beings (with the substantial participation of nanotechnology) may, at first sight, seem limited. Two considerations, however, contest this estimation: first, the vision of the technical enhancement of human beings is actually being seriously advocated. Research projects are being planned in this direction, and milestones for reaching this goal are being set up,[2] whereby nanotechnology takes on the role of an "enabling technology."[18] Furthermore, technical enhancements are by no means completely new, but are—in part—actually established, as the example of plastic surgery as a technical correction of physical characteristics felt to be imperfections shows, and as is the case in the practice of administering psycho-active substances. It isn't difficult to predict that the possibilities and the realization of technical improvements of human beings will increase; demand is conceivable. In view of the moral questions connected with this development and of their conflict potential, ethical reflection is needed in this field.

4. IS AN INDEPENDENT NANO-ETHICS NECESSARY?

Nonetheless, we don't have to reckon with a "Nano-ethics" as a new branch of applied ethics. The propagation of nano-ethics overlooks the fact that many of the ethical questions raised by nanotechnology are already known from other contexts of ethical reflection. The ethics of technology, bioethics, the ethics of medicine or also the theoretical philosophy of technology concern themselves with questions of sustainability, of risk assessment, of the interface between human beings and technology, especially between living beings and technology. These questions are in themselves not new,[23] as the analyses of the individual formulations of the problematic definitions have shown.

Partially new, however, is their convergence in nanotechnology. Analogous to the well-known fact that nanosciences and nanotechnology are fields in which the traditional borders between physics, chemistry, biology, and the engineering sciences are crossed, various traditional lines of ethical reflection also converge in ethical questions in nanotechnology. The fashionable creativity in coining terms, as it precipitates itself in designations like "Neuro-philosophy" or "Nano-ethics," obscures the integrative and cross-sectional nature of many ethical challenges, rather than being particularly helpful. We don't need any new sub-discipline of applied ethics called "nano-ethics" but because new topics and questions are concentrated in nanotechnology[2] and because it accelerates scientific and technical progress, there is a need for ethics in and for nanotechnology. Presupposition is, in particular, willingness on the part of the ethicists for an open reflection on ethical aspects of nanotechnology beyond the scope of classical "hyphenated ethics" and to a dialogue with natural and engineering scientists.

5. ETHICS AS CONCOMITANT REFLECTION OF NANOTECHNOLOGY

Ethics often seems to lag seriously behind technical progress and to fall short of the occasionally great expectations. The rapid pace of

innovation in technicization has the effect that ethical deliberations often come too late; after all of the relevant decisions have already been made, when it is far too late to influence the course of technology development. Technological and scientific progress shapes reality which can no longer be revised after certain points of no return.[22] Ethics in this perspective, could, at best, act as a repair service for problems which have already arisen.

The above is, however, a one-sided view. Ethics actually can provide orientation in the early phases of innovation, e.g., because even the scientific and technical basis includes certain risks which are ethically unacceptable. The opinion that one has to wait with ethical reflection until the corresponding products are on the market and have already caused problems can easily be refuted. This is because the technical knowledge and capabilities are known, as a rule, long before market entry, and can, with the reservation of the well-known problems of the uncertainty of predicting the future, be judged as to their consequences and normative implications. Due to nanotechnology's early stage of development (many prefer to call it nanoscience), we have here a rare case of an advantageous opportunity: there is the chance and also the time for concomitant reflection, as well as the opportunity to integrate the results of reflection into the process of technology design, and thereby to contribute to the further development of nanotechnology.[3]

These considerations don't necessarily mean that ethical deliberations have to be made for absolutely every scientific or technical idea. The problems of a timely occupation with new technologies appear most vividly in the diverse questions raised by the visions of salvation and horror as regards nanotechnology. What sense is there in concerning oneself hypothetically with the ethical aspects of an extreme lengthening of the human life-span,[3] or with self-replicating nanorobots? Most scientists are of the opinion that these are speculations which stem much rather from the realm of science fiction than from problem analysis which is to be taken seriously. We shouldn't forget that ethical reflection binds resources, and there should therefore be a certain evidence for the realizability of these visions, if resources are to be invested in them which could then be lacking elsewhere. Ethical reflection is not necessary "in advance," nor just for the sake of the intellectual diversion it provides. In this respect, we need our own "vision assessment."[24]

Ethical reflection and policy advice built on it are neither senseless nor premature in early phases of development if there are realistic possibilities for the practicability of the technology concerned. Then, there are possibilities for designing ethical assessments as concomitant processes of technological development. If, at first, only rather abstract considerations on the lines of technological development are possible, valuable advice for the further path of development can nonetheless already be given (e.g., by means of timely allusions to potential technology conflicts and to methods of de-escalation). Further, ethical judgment makes orientation for planning the *process* of technological development possible (for example, with regard to questions of equity). In the course of the continuing development of possibilities for the application of nanotechnology, it is then possible to continuously develop the – initially abstract – estimations and orientations on the basis of newly acquired knowledge and, finally, to carry out an ethically-based technology assessment: "Nanoethics is not something one can complete satisfactorily either first or last but something that needs to be continually updated."[3]

The added value in comparison with a later start for ethical reflection is obvious; even the process of technological development profits from the assessment of the consequences, and *ethics* avoids "coming too late," which it is occasionally accused of doing.[3] Technology assessment and ethics have the responsibility, in view of the rapid and momentous developments in nanotechnology, to make the societal process of learning which is always connected with the introduction of a new technology as constructive, transparent, and effective as possible by means of timely investigation and reflection.

REFERENCES

1. Paschen, H., Coenen, C., Fleischer, T., Grünwald, R., Oertel, D., Revermann, C. (2004) Nanotechnologi,. Springer, Berlin et al.
2. Roco, M.C., Bainbridge, W.S. (2002) eds. Converging Technologies for Improving Human Performance. National Science Foundation, Arlington, Virginia.
3. Moor, J., Weckert, J. (2003) Nanoethics: Assessing the Nanoscale from an Ethical Point of View. Lecture, Technical University of Darmstadt, October 10, 2003.
4. Robison, W.L. (2004) Nano-Ethics. Rochester Institute of Technology (manuscript).
5. Mnyusiwalla, A. et al. (2003) Mind the Gap. Science and Ethics in Nanotechnology. Nanotechnology 14: R9–R13; www.iop.org/EJ.
6. Poel, I. van de (2001) Investigating Ethical Issues in Engineering Design. Science and Engineering Ethics 7: 429–446.
7. Grunwald, A. (2000) Against Over-Estimating the Role of Ethics in Technology. Science and Engineering Ethics 6: 181–196.
8. Grunwald, A. (2003) Methodical Reconstruction of Ethical Advises. In: G. Bechmann, I. Hronszky (eds.): Expertise and Its Interfaces. Edition Sigma, Berlin: 103–124.
9. Weckert, J. (2002) Lilliputian Computer Ethics. Metaphilosophy 33: 366–375.
10. Hermerén, G. (2004) Nano Ethics Primer. Conference "Mapping out nano risks." Brussels March 1, 2004.
11. Baumgartner, C. (2004) Ethische Aspekte nanotechnologischer Forschung und Entwicklung in der Medizin. Das Parlament B23-24: 39–46.
12. Jonas, H. (1979) Das Prinzip Verantwortung. Versuch einer Ethik für die technologische Zivilisation. Suhrkamp, Frankfurt.
13. ETC-Group (2003) The Big Down. Atomtech: Technologies Converging at the Nanoscale. www.etcgroup.org.
14. Harremoes, P., Gee, D., MacGarvin, M., Stirling, A., Keys, J., Wynne, B., Guedes Vaz, S. (2002) eds.: The Precautionary Principle in the 20th Century. Late Lessons from Early Warnings. Sage, London.
15. Schomberg, R.V. (2004) The Normative Dimensions of the Precautionary Principle and its Relations to Science and Risk Management Decisions. In: Achen, T. (ed.): Microscopic Modification and Big Politics. Linkoeping Studies in Arts and Sciences, Vadstena (in press).
16. Gethmann, C.F., Mittelstraß, J. (1992) Umweltstandards. GAIA 1: 16–25.
17. Kopfmueller, J., Brandl, V., Jörissen, J., Paetau, M., Banse, G., Coenen, R., Grunwald, A., (2001) Nachhaltige Entwicklung integrativ betrachtet. Edition Sigma, Berlin.
18. Fleischer, T. (2003) Technikgestaltung für mehr Nachhaltigkeit: Nanotechnologie. In: Coenen, R., Grunwald, A. (eds.): Nachhaltigkeitsprobleme in Deutschland. Analyse und Lösungsstrategien. Edition Sigma, Berlin: 356–373.
19. Fleischer, T., Grunwald, A. (2002) Technikgestaltung für mehr Nachhaltigkeit – Anforderungen an die Technikfolgenabschätzung. In: Grunwald, A. (2002, ed.): Technikgestaltung für eine nachhaltige Entwicklung. Von der Konzeption zur Umsetzung. Edition Sigma, Berlin: 95–146.
20. Nanoforum (2004) ed. Nanotechnology. Benefits, Risks, Ethical, legal, and social Aspects of NANOTECHNOLOGY, www.nanoforum.org.
21. Altmann, J., Gubrud, A.A. (2002) Risks from military uses of nanotechnology–the need for Technology Assessment and Preventive Control. In: Roco, M., Tonellini, R. (eds.): Nanotechnology—Revolutionary Opportunities and Societal Implications, European Commission, Luxembourg.
22. Habermas, J. (2001) Die Zukunft der menschlichen Natur, Suhrkamp, Frankfurt.
23. MacDonald, C. (2004) Nanotech is Novel; the Ethical Issues Are Not. The Scientist 18: 3.
24. Grunwald, A. (2004) Vision Assessment as a new element of the Technology Futures Analyses toolbox. EU-US Seminar: New Technology Foresight, Forecasting & Assessment Methods-Seville 13–14 May 2004.

2.4 INTERNET AND SOCIAL MEDIA

2.4.1 IS GOOGLE MAKING US STUPID?

NICHOLAS CARR

In this selection, Nicholas Carr argues that easy access to information, especially via the Internet, is fundamentally changing the ways we learn, and therefore think. Drawing on examples as diverse as HAL, the rogue computer in *2001: A Space Odyssey* and comments by Google entrepreneurs Sergey Brin and Larry Page, Carr suggests that the way we acquire and process written information today has changed drastically since the Internet came into our lives.

There is no doubt that the Internet provides us with endless new opportunities to locate information in comparison to the research methods of the past—and the quantity of material we read today may well be more than people ever have. But it is in the very nature of these endless options to serve as distractions that, according to Carr, interfere with deep thought processes and learning. As readers, many of us no longer have the ability to become deeply absorbed in books; our attention wanders and we are easily distracted, a phenomenon that Carr suggests may be due to changes produced by an information diet consisting largely of online digital media. Online reading may have many advantages, but it may also be costing us the ability to read and to think deeply and with sustained concentration.

Critics of earlier information technologies, such as writing and the printing press, warned of the adverse effects of their widespread adoption. Like them, Carr suggests that in evaluating the social impact of the Internet and new digital information technologies, we should look carefully at both the benefits and the costs.

FOCUS QUESTIONS

1. How, according to Carr, have recent information technologies such as Google have changed us, compared with other types of technologies? What do you feel might be the long-range impact of such modifications, if the influence of such technologies continues to grow?
2. What has been your experiences surfing the Web, and do you feel that it has benefited your education? Since it is likely that you grew up having the Internet and Google always available to you, is it possible for you to even imagine what life was like before these services? Do your best to explain the significance of the shift.
3. The desire to reduce costs and increase productivity has been a major goal of capitalist societies for centuries. In what ways do you feel Google and other information technologies have helped those efforts, and what might be some negative effects of that process, as discussed in this selection? Explain why you believe ethical concerns over such productivity shifts are valid or not.

KEYWORDS

algorithm, alphabet, bounce, cognition, deep reading, ideogram, intellectual technologies, information-processing device, knowledge work, malleability, neural circuitry, skimming activity, Taylorism, universal medium, utilitarian resource

"Dave, stop. Stop, will you? Stop, Dave. Will you stop, Dave?" So the supercomputer HAL pleads with the implacable astronaut Dave Bowman in a famous and weirdly poignant scene toward the end of Stanley Kubrick's *2001: A Space Odyssey*. Bowman, having nearly been sent to a deep-space death by the malfunctioning machine, is calmly, coldly disconnecting the memory circuits that control its artificial brain. "Dave, my mind is going," HAL says, forlornly. "I can feel it. I can feel it."

I can feel it, too. Over the past few years I've had an uncomfortable sense that someone, or something, has been tinkering with my brain, remapping the neural circuitry, reprogramming the memory. My mind isn't going—so far as I can tell—but it's changing. I'm not thinking the way I used to think. I can feel it most strongly when I'm reading. Immersing myself in a book or a lengthy article used to be easy. My mind would get caught up in the narrative or the turns of the argument, and I'd spend hours strolling through long stretches of prose. That's rarely the case anymore. Now my concentration often starts to drift after two or three pages. I get fidgety, lose the thread, begin looking for something else to do. I feel as if I'm always dragging my wayward brain back to the text. The deep reading that used to come naturally has become a struggle.

I think I know what's going on. For more than a decade now, I've been spending a lot of time online, searching and surfing and sometimes adding to the great databases of the Internet. The Web has been a godsend to me as a writer. Research that once required days in the stacks or periodical rooms of libraries can now be done in minutes. A few Google searches, some quick clicks on hyperlinks, and I've got the telltale fact or pithy quote I was after. Even when I'm not working, I'm as likely as not to be foraging in the Web's info-thickets—reading and writing e-mails, scanning headlines and blog posts, watching videos and listening to podcasts, or just tripping from link to link to link. (Unlike footnotes, to which they're sometimes likened, hyperlinks don't merely point to related works; they propel you toward them.)

For me, as for others, the Net is becoming a universal medium, the conduit for most of the information that flows through my eyes and ears and into my mind. The advantages of having immediate access to such an incredibly rich store of information are many, and they've been widely described and duly applauded. "The perfect recall of silicon memory," *Wired*'s Clive Thompson has written, "can be an enormous boon to thinking." But that boon comes at a price. As the media theorist Marshall McLuhan pointed out in the 1960s, media are not just passive channels of information. They supply the stuff of thought, but they also shape the process of thought. And what the Net seems to be doing is chipping away my capacity for concentration and contemplation. My mind now expects to take in information the way the Net distributes it: in a swiftly moving stream of particles. Once I was a scuba diver in the sea of words. Now I zip along the surface like a guy on a Jet Ski.

I'm not the only one. When I mention my troubles with reading to friends and acquaintances—literary types, most of them—many say they're having similar experiences. The more they use the Web, the more they have to fight to stay focused on long pieces of writing. Some of the bloggers I follow have also begun mentioning the phenomenon. Scott Karp, who writes a blog about online media, recently confessed that he has stopped reading books altogether. "I was a lit major in college, and used to be [a] voracious book reader," he wrote. "What happened?" He speculates on the answer: "What if I do all my reading on the web not so much because the way I read has changed, i.e. I'm just seeking convenience, but because the way I THINK has changed?"

Bruce Friedman, who blogs regularly about the use of computers in medicine, also has described how the Internet has altered his mental habits. "I now have almost totally lost the ability to read and absorb a longish article on the web or in print," he wrote earlier this year. A pathologist who has long been on the faculty of the University of Michigan Medical School, Friedman elaborated on his comment in a telephone conversation with me. His thinking, he said, has taken on a "staccato" quality, reflecting the way he quickly scans short passages of text from many sources online. "I can't read *War and Peace* anymore," he admitted. "I've lost the ability to do that. Even a blog post of more than three or four paragraphs is too much to absorb. I skim it."

Anecdotes alone don't prove much. And we still await the long-term neurological and psychological experiments that will provide a definitive picture of how Internet use affects cognition. But a recently published study of online research habits, conducted by scholars from University College London, suggests that we may well be in the midst of a sea change in the way we read and think. As part of the five-year research program, the scholars examined computer logs documenting the behavior of visitors to two popular research sites, one operated by the British Library and one by a U.K. educational consortium, that provide access to journal articles, e-books, and other sources of written information. They found that people using the sites exhibited "a form of skimming activity," hopping from one source to another and rarely returning to any source they'd already visited. They typically read no more than one or two pages of an article or book before they would "bounce" out to another site. Sometimes they'd save a long article, but there's no evidence that they ever went back and actually read it. The authors of the study report:

> It is clear that users are not reading online in the traditional sense; indeed there are signs that new forms of "reading" are emerging as users "power browse" horizontally through titles, contents pages and abstracts going for quick wins. It almost seems that they go online to avoid reading in the traditional sense.

Thanks to the ubiquity of text on the Internet, not to mention the popularity of text-messaging on cell phones, we may well be reading more today than we did in the 1970s or 1980s, when television was our medium of choice. But it's a different kind of reading, and behind it lies a different kind of thinking—perhaps even a new sense of the self. "We are not only *what* we read," says Maryanne Wolf, a developmental psychologist at Tufts University and the author of *Proust and the Squid: The Story and Science of the Reading Brain*. "We are *how* we read." Wolf worries that the style of reading promoted by the Net, a style that puts "efficiency" and "immediacy" above all else, maybe weakening our capacity for the kind of deep reading that emerged when an earlier technology, the printing press, made long and complex works of prose commonplace. When we read online, she says, we tend to become "mere decoders of information." Our ability to interpret text, to make the rich mental connections that form when we read deeply and without distraction, remains largely disengaged.

Reading, explains Wolf, is not an instinctive skill for human beings. It's not etched into our genes the way speech is. We have to teach our minds how to translate the symbolic characters we see into the language we understand. And the media or other technologies we use in learning and practicing the craft of reading play an important part in shaping the neural circuits inside our brains. Experiments demonstrate that readers of ideograms, such as the Chinese, develop a mental circuitry for reading that is very different from the circuitry found in those of us whose written language employs an alphabet. The variations extend across many regions of the brain, including those that govern such essential cognitive functions as memory and the interpretation of visual and auditory stimuli. We can expect as well that the circuits woven by our use of the Net will be different from those woven by our reading of books and other printed works.

Sometime in 1882, Friedrich Nietzsche bought a typewriter—a Malling-Hansen Writing Ball, to be precise. His vision was failing, and keeping his eyes focused on a page had become exhausting and painful, often bringing on crushing

headaches. He had been forced to curtail his writing, and he feared that he would soon have to give it up. The typewriter rescued him, at least for a time. Once he had mastered touch-typing, he was able to write with his eyes closed, using only the tips of his fingers. Words could once again flow from his mind to the page.

But the machine had a subtler effect on his work. One of Nietzsche's friends, a composer, noticed a change in the style of his writing. His already terse prose had become even tighter, more telegraphic. "Perhaps you will through this instrument even take to a new idiom," the friend wrote in a letter, noting that, in his own work, his "'thoughts' in music and language often depend on the quality of pen and paper."

"You are right" Nietzsche replied, "our writing equipment takes part in the forming of our thoughts." Under the sway of the machine, writes the German media scholar Friedrich A. Kittler, Nietzsche's prose "changed from arguments to aphorisms, from thoughts to puns, from rhetoric to telegram style."

The human brain is almost infinitely malleable. People used to think that our mental meshwork, the dense connections formed among the 100 billion or so neurons inside our skulls, was largely fixed by the time we reached adulthood. But brain researchers have discovered that that's not the case. James Olds, a professor of neuroscience who directs the Krasnow Institute for Advanced Study at George Mason University, says that even the adult mind "is very plastic." Nerve cells routinely break old connections and form new ones. "The brain," according to Olds, "has the ability to reprogram itself on the fly, altering the way it functions."

As we use what the sociologist Daniel Bell has called our "intellectual technologies"—the tools that extend our mental rather than our physical capacities—we inevitably begin to take on the qualities of those technologies. The mechanical clock, which came into common use in the 14th century, provides a compelling example. In *Technics and Civilization*, the historian and cultural critic Lewis Mumford described how the clock "disassociated time from human events and helped create the belief in an independent world of mathematically measurable sequences." The "abstract framework of divided time" became "the point of reference for both action and thought."

The clock's methodical ticking helped bring into being the scientific mind and the scientific man. But it also took something away. As the late MIT computer scientist Joseph Weizenbaum observed in his 1976 book, *Computer Power and Human Reason: From Judgment to Calculation*, the conception of the world that emerged from the widespread use of timekeeping instruments "remains an impoverished version of the older one, for it rests on a rejection of those direct experiences that formed the basis for, and indeed constituted, the old reality." In deciding when to eat, to work, to sleep, to rise, we stopped listening to our senses and started obeying the clock.

The process of adapting to new intellectual technologies is reflected in the changing metaphors we use to explain ourselves to ourselves. When the mechanical clock arrived, people began thinking of their brains as operating "like clockwork." Today, in the age of software, we have come to think of them as operating "like computers." But the changes, neuroscience tells us, go much deeper than metaphor. Thanks to our brains plasticity, the adaptation occurs also at a biological level.

The Internet promises to have particularly far-reaching effects on cognition. In a paper published in 1936, the British mathematician Alan Turing proved that a digital computer, which at the time existed only as a theoretical machine, could be programmed to perform the function of any other information-processing device. And that's what we're seeing today. The Internet, an immeasurably powerful computing system, is subsuming most of our other intellectual technologies. It's becoming our map and our clock, our printing press and our typewriter, our calculator and our telephone, and our radio and TV.

When the Net absorbs a medium, that medium is recreated in the Net's image. It injects the mediums content with hyperlinks, blinking ads, and other digital gewgaws, and it surrounds

the content with the content of all the other media it has absorbed. A new e-mail message, for instance, may announce its arrival as we're glancing over the latest headlines at a newspaper's site. The result is to scatter our attention and diffuse our concentration.

The Net's influence doesn't end at the edges of a computer screen, either. As people's minds become attuned to the crazy quilt of Internet media, traditional media have to adapt to the audience's new expectations. Television programs add text crawls and pop-up ads, and magazines and newspapers shorten their articles, introduce capsule summaries, and crowd their pages with easy-to-browse info-snippets. When, in March of this year, *The New York Times* decided to devote the second and third pages of every edition to article abstracts, its design director, Tom Bodkin, explained that the "shortcuts" would give harried readers a quick "taste" of the day's news, sparing them the "less efficient" method of actually turning the pages and reading the articles. Old media have little choice but to play by the new-media rules.

Never has a communications system played so many roles in our lives—or exerted such broad influence over our thoughts—as the Internet does today. Yet, for all that's been written about the Net, there's been little consideration of how, exactly, it's reprogramming us. The Net's intellectual ethic remains obscure.

About the same time that Nietzsche started using his typewriter, an earnest young man named Frederick Winslow Taylor carried a stopwatch into the Midvale Steel plant in Philadelphia and began a historic series of experiments aimed at improving the efficiency of the plant's machinists. With the approval of Midvale's owners, he recruited a group of factory hands, set them to work on various metalworking machines, and recorded and timed their every movement as well as the operations of the machines. By breaking down every job into a sequence of small, discrete steps and then testing different ways of performing each one, Taylor created a set of precise instructions—an "algorithm," we might say today—for how each worker should work. Midvale's employees grumbled about the strict new regime, claiming that it turned them into little more than automatons, but the factory's productivity soared.

More than a hundred years after the invention of the steam engine, the Industrial Revolution had at last found its philosophy and its philosopher. Taylor's tight industrial choreography—his "system," as he liked to call it—was embraced by manufacturers throughout the country and, in time, around the world. Seeking maximum speed, maximum efficiency, and maximum output, factory owners used time-and-motion studies to organize their work and configure the jobs of their workers. The goal, as Taylor defined it in his celebrated 1911 treatise, *The Principles of Scientific Management*, was to identify and adopt, for every job, the "one best method" of work and thereby to effect "the gradual substitution of science for rule of thumb throughout the mechanical arts." Once his system was applied to all acts of manual labor, Taylor assured his followers, it would bring about a restructuring not only of industry but of society, creating a utopia of perfect efficiency. "In the past the man has been first," he declared; "in the future the system must be first."

Taylor's system is still very much with us; it remains the ethic of industrial manufacturing. And now, thanks to the growing power that computer engineers and software coders wield over our intellectual lives, Taylor's ethic is beginning to govern the realm of the mind as well. The Internet is a machine designed for the efficient and automated collection, transmission, and manipulation of information, and its legions of programmers are intent on finding the "one best method"—the perfect algorithm—to carry out every mental movement of what we've come to describe as "knowledge work."

Google's headquarters, in Mountain View, California—the Googleplex—is the Internet's high church, and the religion practiced inside its walls is Taylorism. Google, says its chief executive, Eric Schmidt, is "a company that's founded around the science of measurement," and it is striving to "systematize everything" it does. Drawing on the terabytes of behavioral

data it collects through its search engine and other sites, it carries out thousands of experiments a day, according to the *Harvard Business Review*, and it uses the results to refine the algorithms that increasingly control how people find information and extract meaning from it. What Taylor did for the work of the hand, Google is doing for the work of the mind.

The company has declared that its mission is "to organize the world's information and make it universally accessible and useful." It seeks to develop "the perfect search engine," which it defines as something that "understands exactly what you mean and gives you back exactly what you want." In Google's view, information is a kind of commodity, a utilitarian resource that can be mined and processed with industrial efficiency. The more pieces of information we can "access" and the faster we can extract their gist, the more productive we become as thinkers.

Where does it end? Sergey Brin and Larry Page, the gifted young men who founded Google while pursuing doctoral degrees in computer science at Stanford, speak frequently of their desire to turn their search engine into an artificial intelligence, a HAL-like machine that might be connected directly to our brains. "The ultimate search engine is something as smart as people—or smarter," Page said in a speech a few years back. "For us, working on search is a way to work on artificial intelligence." In a 2004 interview with *Newsweek*, Brin said, "Certainly if you had all the world's information directly attached to your brain, or an artificial brain that was smarter than your brain, you'd be better off." Last year, Page told a convention of scientists that Google is "really trying to build artificial intelligence and to do it on a large scale."

Such an ambition is a natural one, even an admirable one, for a pair of math whizzes with vast quantities of cash at their disposal and a small army of computer scientists in their employ. A fundamentally scientific enterprise, Google is motivated by a desire to use technology, in Eric Schmidt's words, "to solve problems that have never been solved before," and artificial intelligence is the hardest problem out there. Why wouldn't Brin and Page want to be the ones to crack it?

Still, their easy assumption that we'd all "be better off" if our brains were supplemented, or even replaced, by an artificial intelligence is unsettling. It suggests a belief that intelligence is the output of a mechanical process, a series of discrete steps that can be isolated, measured, and optimized. In Google's world, the world we enter when we go online, there's little place for the fuzziness of contemplation. Ambiguity is not an opening for insight but a bug to be fixed. The human brain is just an outdated computer that needs a faster processor and a bigger hard drive.

The idea that our minds should operate as high-speed data-processing machines is not only built into the workings of the Internet, it is the network's reigning business model as well. The faster we surf across the Web—the more links we click and pages we view—the more opportunities Google and other companies gain to collect information about us and to feed us advertisements. Most of the proprietors of the commercial Internet have a financial stake in collecting the crumbs of data we leave behind as we flit from link to link—the more crumbs, the better. The last thing these companies want is to encourage leisurely reading or slow, concentrated thought. It's in their economic interest to drive us to distraction.

Maybe I'm just a worrywart. Just as there's a tendency to glorify technological progress, there's a countertendency to expect the worst of every new tool or machine. In Plato's *Phaedrus*, Socrates bemoaned the development of writing. He feared that, as people came to rely on the written word as a substitute for the knowledge they used to carry inside their heads, they would, in the words of one of the dialog's characters, "cease to exercise their memory and become forgetful." And because they would be able to "receive a quantity of information without proper instruction," they would "be thought very knowledgeable when they are for the most part quite ignorant." They would be "filled with the conceit of

wisdom instead of real wisdom." Socrates wasn't wrong—the new technology did often have the effects he feared—but he was shortsighted. He couldn't foresee the many ways that writing and reading would serve to spread information, spur fresh ideas, and expand human knowledge (if not wisdom).

The arrival of Gutenberg's printing press, in the 15th century, set off another round of teeth gnashing. The Italian humanist Hieronimo Squarciafico worried that the easy availability of books would lead to intellectual laziness, making men "less studious" and weakening their minds. Others argued that cheaply printed books and broadsheets would undermine religious authority, demean the work of scholars and scribes, and spread sedition and debauchery. As New York University professor Clay Shirky notes, "Most of the arguments made against the printing press were correct, even prescient." But, again, the doomsayers were unable to imagine the myriad blessings that the printed word would deliver.

So, yes, you should be skeptical of my skepticism. Perhaps those who dismiss critics of the Internet as Luddites or nostalgists will be proved correct, and from our hyperactive, data-stoked minds will spring a golden age of intellectual discovery and universal wisdom. Then again, the Net isn't the alphabet, and although it may replace the printing press, it produces something altogether different. The kind of deep reading that a sequence of printed pages promotes is valuable not just for the knowledge we acquire from the author's words but for the intellectual vibrations those words set off within our own minds. In the quiet spaces opened up by the sustained, undistracted reading of a book, or by any other act of contemplation, for that matter, we make our own associations, draw our own inferences and analogies, foster our own ideas. Deep reading, as Maryanne Wolf argues, is indistinguishable from deep thinking.

If we lose those quiet spaces, or fill them up with "content," we will sacrifice something important not only in our selves but in our culture. In a recent essay, the playwright Richard Foreman eloquently described what's at stake:

> I come from a tradition of Western culture, in which the ideal (my ideal) was the complex, dense and "cathedral-like" structure of the highly educated and articulate personality—a man or woman who carried inside themselves a personally constructed and unique version of the entire heritage of the West. [But now] I see within us all (myself included) the replacement of complex inner density with a new kind of self—evolving under the pressure of information overload and the technology of the "instantly available."

As we are drained of our "inner repertory of dense cultural inheritance," Foreman concluded, we risk turning into "'pancake people'—spread wide and thin as we connect with that vast network of information accessed by the mere touch of a button."

I'm haunted by that scene in *2001*. What makes it so poignant, and so weird, is the computer's emotional response to the disassembly of its mind: its despair as one circuit after another goes dark, its childlike pleading with the astronaut—"I can feel it. I can feel it. I'm afraid"—and its final reversion to what can only be called a state of innocence. HAL's outpouring of feeling contrasts with the emotionlessness that characterizes the human figures in the film, who go about their business with an almost robotic efficiency. Their thoughts and actions feel scripted, as if they're following the steps of an algorithm. In the world of *2001*, people have become so machinelike that the most human character turns out to be a machine. That's the essence of Kubrick's dark prophecy: as we come to rely on computers to mediate our understanding of the world, it is our own intelligence that flattens into artificial intelligence.

2.4.2 CORPORATE CENSORSHIP

REBECCA MACKINNON

The Internet is a primary new arena in the struggle for freedom of thought and expression, according to Rebecca MacKinnon. MacKinnon is a former CNN journalist, who headed its bureaus in Beijing and Tokyo. She is also a founder of Global Voices Online, an Internet site that hosts blogs from around the world. She has been a fellow at the New America Foundation, and has spoken widely on the issue of Internet freedom.

This reading from MacKinnon's book, *Consent of the Networked: The Worldwide Struggle for Internet Freedom,* examines how private companies, such as Apple, Google, and Verizon, are attempting to control the kinds of political content that users can post using their Web services. In an interview she gave when this book was published, MacKinnon said, "Governments and corporations are working actively to shape the Internet to fit their own needs. The most insidious situations arise when both government and corporations combine their efforts to exercise power over the same people at the same time, in largely unconstrained and unaccountable ways. This is why I argue that if we the people do not wake up and fight for the protection of our own rights and interests on the Internet, we should not be surprised to wake up one day to find that they have been programmed, legislated, and sold away."

FOCUS QUESTIONS

1. MacKinnon provides several examples of cases in which private Internet companies have censored the free expression of ideas or information. In some cases, these practices are due to laws imposed by governments, but in other cases they are a result of the companies' own policies. Is one kind of censorship more problematic than the other? Explain.
2. What is *net neutrality*? When Tim Wu coined this term, he argued that the history of communications technologies follows a pattern of initially open access that becomes increasingly restricted and controlled. Do you see evidence for this happening today with the Internet? Explain.
3. What are the main arguments for and against unfettered Internet freedom? Why might both of these "purist" positions be inadequate? Explain.
4. Should there be different approaches to Internet freedom as it applies to wired (cable, DSL, etc.) connections and wireless Internet communications, such as smartphones? Why might some companies wish to exert greater control over wireless content? Explain.
5. What are some examples of "app censorship?" Do you believe that companies such as Apple and YouTube should be able to control the content that users post on their sites, based on political, religious, or moral grounds? Explain.

Source: From *Consent of the Networked: The Worldwide Struggle for Internet Freedom* by Rebecca MacKinnon, p. 115–129. Copyright © 2012 Rebecca MacKinnon. Reprinted by permission of Basic Books, a member of the Perseus Book Group.

KEYWORDS

app censorship, Apple, Facebook, Google, Internet freedom, ISPs, mobile phones, net neutrality, Verizon, wireless Internet technology

In the fall of 2009, Apple officially launched the iPhone in China in partnership with a domestic mobile carrier, China Unicom. As a condition for entry into the Chinese market, Apple had to agree to the Chinese government's censorship criteria in vetting the content of all iPhone applications—or "apps"—available for download on devices sold in mainland China. (Most apps are created by independent developers—individuals, companies, or organizations—then submitted to Apple for approval and inclusion in its app store.) On Apple's special store for the Chinese market, apps related to the Dalai Lama are censored, as is one containing information about the exiled Uighur dissident leader Rebiya Kadeer. Apple similarly censors apps for iPads sold in China. So much for that revolutionary, Big Brother-destroying 1984 Super Bowl ad. Fifteen years later, Apple seems quite willing to accommodate Big Brother's demands for the sake of market access.

Internet and mobile telecommunications companies (whose functions and relationships are increasingly intertwined) create computer code that functions as a kind of law, in that it shapes what people can do and sometimes directly censors what they can see. Sometimes the computer code is written to comply with government demands, as in the Chinese case. Sometimes government uses legislative code—law—to force the evolution of computer code in one direction or the other. Sometimes these laws reflect the "consent of the governed," and sometimes they do not, depending on what sort of government is doing the lawmaking. Sometimes the laws are arguably more reflective of certain powerful and wealthy lobbies than the public as a whole. Sometimes the law intends to reflect the wider public interest but actually reflects politicians' ignorance of technology thus creating unintended and sometimes negative consequences. At any rate, legislative code does not cover most aspects of computer code, and it certainly fails to anticipate innovations, as Lawrence Lessig first pointed out a decade ago. Neither citizens nor businesses have an interest in excessive micromanagement of computer code by legislative code—or innovation would be impossible. Thus, much of the time companies apply their own criteria, based mainly on commercial factors. Unfortunately, the impact of decisions on customers' ability to exercise their political rights and freedoms is at best an afterthought, if the question even comes to managers' and engineers' minds at all.

NET NEUTRALITY

In 2007, the abortion rights group NARAL Pro-Choice America sent out text messages to supporters who had signed up to receive them. One message said, "End Bush's global gag rule against birth control for world's poorest women! Call Congress. (202) XXX-XXXX.Naral Text4-Choice." Verizon Wireless blocked the NARAL messages.

When challenged, company executives said it was company policy to block messages from any group "that seeks to promote an agenda or distribute content that, in [Verizon's] discretion, may be seen as controversial or unsavory." Then, after further news coverage and public outcry, a spokesperson said that the initial explanation had been "an incorrect interpretation of a dusty internal policy … designed to ward against communications such as anonymous hate messaging and adult materials sent to children." Verizon did not actually need to explain itself, however: standard customer agreements generally include language stating that the customer understands that service can be cut off and data transmissions blocked for any reason.

In late 2010,T-Mobile blocked a medical marijuana dispensary listing service based in California (where medical marijuana is legal) from using its short code text messaging service, claiming that messages had been sent on behalf of a group of which T-Mobile did not approve.

Though the group sued T-Mobile, the company called the lawsuit "without merit." (The case was dismissed by the Southern District Court "with prejudice" against the defendant, T-Mobile.) Indeed, although US law prohibits phone companies from blocking calls placed on their networks, the law does not apply to data and text messages, so as far as the law is concerned, companies are free to do as they please.

Free speech groups point to other troubling examples of political censorship by companies: In 2007 during a live-streaming broadcast of a Pearl Jam concert, AT&T muted the sound when lead singer Eddie Vedder made comments critical of then-president George W. Bush. Comcast has been accused of blocking entire peer-to-peer applications because some people were using them to share copyrighted works illegally, thus hindering the legal use of these networks by people trying to share large video files, which inevitably include homemade videos containing political speech. Granted, in the United States people have other ways of expressing themselves and getting information out. But the fact that citizen efforts to conduct legal political activities over private networks can be arbitrarily stymied by private service operators is an example of how the political discourse in a democracy can be eroded if citizens are not vigilant about exposing political discrimination when it occurs and holding companies publicly accountable.

There is further concern that even when companies do not discriminate against specific content, they can still discriminate for and against different categories of Internet traffic. This ability could potentially make it more difficult for independent and nonprofit citizen media to reach large audiences or build broad communities. If ISPs are allowed to discriminate, they can also block other content services that happen to compete with their own.

Internet service providers now have the technical ability to "see" what kinds of application a particular subscriber is using at a given moment in time: Are you using Skype? YouTube? Streaming movies on Netflix? Or are you accessing a nonprofit website that uses the open-source WordPress platform? Your service provider knows. Currently in most countries, there is no law preventing Internet service providers from discriminating between services for a profit: your ISP could in theory offer a "tiered" access package in which access to certain services belonging to major brand names (Amazon, Netflix, YouTube, and Facebook, for example) would be cheaper than access to the general Internet or to lesser known applications. Brand-name companies could pay the ISP for the privilege of gaining preferential access to users at lower or no cost. Or an ISP could demand money from popular services in exchange for smooth and fast delivery.

In late 2010 Comcast attempted to move in this direction when it demanded that Level 3 Communications, a major client of Netflix (which also competes with Comcast's own video services, particularly now that Comcast has acquired a controlling stake in NBC Universal), should pay an extra fee in exchange for smooth delivery of Netflix videos to Comcast broadband subscribers. Level 3 accused Comcast of trying to set up what amounts to a toll booth on the Internet. The Internet's empowering nature has been based largely on the fact that any citizen can potentially create media and reach global audiences just by sending a tweet or publishing a blog. On a "neutral" Internet, it doesn't matter what platform the citizen uses—that person's video, blog post, or podcast is equally accessible to anybody anywhere (as long as a government isn't trying to censor it). On a non-neutral Internet in which platforms and applications must compete by paying fees to gain maximum access to Internet users—and in which Internet users are asked to pay different amounts for different tiers of access determined by the different deals the ISP has cut with the companies that run different services and platforms—citizens will depend all the more on large, well-financed, commercially operated platforms to carry out political discourse and political organizing.

The obvious question, then, is: How do citizens make sure that private agendas and pursuit of profit do not erode consumer choice and even democratic expression?

Some people advocate government regulation as the best solution to this problem. Companies,

they argue, should be banned by law from discriminating against content passing through privately operated networks. Networks should be required by law to be "neutral" in their handling of content—hence the term "net neutrality," coined originally by Tim Wu, a Columbia University law professor. In his influential 2010 book, *The Master Switch: The Rise and Fall of Information Empires*, Wu describes how technologies—from the telegraph and the telephone to radio and film—go through what he calls the "Cycle": new technologies start out open and revolutionary, then eventually get locked down by monopolists in ways that make them less appropriate or even inhospitable to uses that would be considered excessively political, edgy, provocative, or controversial. Yet Wu is also concerned about government interference in private networks in ways that violate citizen rights. He cites AT&T's role in the NSA's warrantless wiretapping as a clear example of how corporate power can be abused by an overreaching government.

Such abuse is most extreme, and indeed commonplace, in authoritarian countries where networks are by definition non-neutral. It has become standard practice for authoritarian regimes to require companies operating digital networks and platforms to censor and discriminate in an opaque and unaccountable manner. Furthermore, although networks and platforms carry out censorship at government direction, the lack of transparency about what is being censored and manipulated— and at whose behest—also makes it very easy for companies to discriminate in any number of ways that serve the interests of both the company and its executives. Citizens are thus vulnerable to abuse of their rights to speech and assembly not only from government but also from private actors. In democracies, it follows that citizens must guard against violation of their digital rights by governments and corporations—or both acting in concert—regardless of whether the company involved is censoring and discriminating on its own initiative or acting under pressure from authorities.

In late 2010 the Federal Communications Commission (FCC) issued rules stipulating basic net neutrality standards for the United States, with some caveats that free speech groups have criticized. Net neutrality proponents including Democratic Senator Al Franken believe it is a *prerequisite* for Internet freedom, and ultimately as important as the First Amendment when it comes to ensuring a free and democratic discourse. Proponents tend to trust democratically elected government more than they trust private companies to act in the public interest.

Opponents of net neutrality argue that government interference of any kind with how private companies operate their networks runs counter to "Internet freedom." They also believe that as more and more Internet users turn to video and interactive gaming, network operators should be free to manage their traffic in whatever way is most efficient. They argue that market competition is the most effective way to keep companies honest. They also argue that excessive regulation stifles innovation, which ultimately hurts citizens, who are deprived of new products and services that cannot take root or thrive, thanks to regulations that inevitably fail to anticipate how technology will evolve.

There are certainly many strong reasons to be concerned about government meddling in private networks. But at the same time, market competition keeps companies honest only when there are ample choices and alternatives, and if the cost of switching services is not prohibitively high for most people. In the United States, that condition does not even come close to being met. Most Americans face very limited choices of broadband and wireless providers, especially in comparison to Europe, where regulations force competitive use of infrastructure.

As *The Economist* pointed out in December 2010 after the FCC issued rules on net neutrality that failed to satisfy purists on either side of the debate, the entire argument obscured "the failure in America to tackle the underlying lack of competition in the provision of Internet access." That, of course, would require much heavier political lifting in Congress: "Getting America's phone and cable companies to open up their networks to others would be a lot harder for politicians than prattling on about neutrality; but it would

do far more to open up the net." As of 2011, there is no sign that Congress is serious about tackling this core problem of monopoly and pseudo-monopoly held by many wireless and broadband companies in many parts of the country. It is highly unlikely that any progress will be made in 2012, an election year in which all major telecommunications companies will contribute generously to congressional and presidential campaigns.

Responding to arguments by civic groups that net neutrality is the best way to support and nurture the growth of a robust digital commons, in 2010 Chile became the first nation to enshrine net neutrality into law, requiring that ISPs "ensure access to all types of content, services or applications available on the network and offer a service that does not distinguish content, applications or services, based on the source of it or their property."

In late 2010 the European Parliament and European Commission held a joint net neutrality summit, at which they agreed that it is permissible for companies to interfere with Internet traffic as long as the companies inform consumers that it is happening. But as net neutrality advocates point out, there remains a concern that as companies grow more comfortable with overtly interfering with customer communications in pursuit of their own interests, it will be easier for them to comply with government requests to conduct surveillance and censorship. Thus it was seen as an important step forward in June 2011 when the Netherlands became the first European country to make net neutrality the law when its parliament voted to ban mobile telephone operators from blocking or charging consumers extra money for using Internet-based communications services.

In many countries a lack of net neutrality makes censorship—whether by companies, government, or some mix of the two—much easier to implement and much less publicly visible, let alone accountable. In 2009 Jorge Bossio wrote a report on the Internet in Peru for the nonprofit Association for Progressive Communications in which he pointed out that Telefonica del Peru had a clause in the customer agreement for its "Speedy home Internet service" that stipulated that "the customer is obliged to make use of content in accordance with the law, the present General Conditions, generally accepted moral values and good manners, and the public order," and that customers could be penalized for transmitting or posting content that was "false, ambiguous, inaccurate, exaggerated or untimely." In so doing, Bossio points out, the ISP had granted itself "powers that correspond to the Peruvian administrative and judicial authorities."

Lack of neutrality has other consequences around the world, which people in wealthy industrialized nations often fail to consider. Technologists in the developing world point out that if networks in the world's major Internet markets cease to be neutral, poor and developing nations will be at an even greater disadvantage in the global marketplace of online content and services. As the debates intensified in the United States and Europe in late 2010, Kenyan engineer and blogger Tom Makau wrote a post outlining what the outcome could mean for his continent: "With the lack of net neutrality, Africa-originated traffic will not be treated equally on the internet playing field as say traffic from Google or Microsoft," he wrote. "This is because the latter will be paying millions of dollars to have their traffic receive preferential treatment from the large data carriers. This, coupled with the fact that the continent produces slightly less than 2% of the total world traffic, will mean African content will be nearly invisible on the Internet."

Mobile Complications

Even companies that support basic neutrality requirements for broadband Internet are taking a more relaxed view when it comes to mobile Internet access. Activists in the developing world who are waging an uphill battle for the right to use the Internet as a space for dissent and political mobilization were distressed by news in August 2010 that Google and Verizon had worked out a common position on net neutrality policy. The two companies agreed that "wireline broadband providers" should not be allowed "to discriminate against or prioritize lawful Internet content in

a way that causes harm to users or competition." However, they also agreed that wireless broadband providers should be exempt from such regulation and ought to be free to manage and prioritize their traffic. Several months later, the FCC issued rules that generally reflected this compromise: though mobile carriers would be prohibited from blocking websites, they would be free to block applications or services unless those applications directly competed with providers' voice and video products, such as Skype. The rules allowed for some "network management," which includes prioritization of some services over others, and did not rule out "paid prioritization" of services, with the stipulation that carriers must be transparent about how they implement their traffic management. Net neutrality opponents and proponents were both highly critical with the compromise position.

Many net neutrality advocates are particularly concerned that although there is still hope that the broadband-based Internet in homes and offices can remain relatively neutral at least in democratic countries, policy and industry trends point to the likelihood that the *mobile* Internet *everywhere* will be much more constrained and manipulated, with limited room for the citizen commons. At the same time, there is a major demographic shift under way when it comes to global Internet use. The Internet is shifting from being a medium used primarily by people who live in economically developed, democratic countries to being a medium whose fastest user growth is in poor or developing countries—and in countries that are authoritarian, quasi-authoritarian, or very weak democracies. Much of the Internet's growth is happening through mobile phones, not through laptops or desktops connected to broadband.

It is therefore likely that Internet users around the world could be divided broadly into two classes: In one class are the wealthy, who have home broadband connections and who own PCs or laptops, and who are empowered by a much more open and neutral Internet. In the other class are the poor, the migrants, and the generally disenfranchised, who use the Internet mainly through highly controlled, easily tracked, non-neutral mobile devices. The situation is all the more insidious when we know that many of these companies are responding to government orders to censor and track users, in ways that are neither transparent nor accountable to the people who have come to depend on them.

What might a non-neutral wireless Internet service actually be like for consumers? In the developing world a picture is already emerging. In May 2010, Facebook launched a service called Facebook Zero in more than forty-five wireless markets, signing deals with mobile carriers to provide Facebook access free to subscribers. Now in a growing number of countries including Pakistan, India, Sudan, and Tunisia, mobile phone users must pay the standard data rate to browse the open Web but can access the special free Facebook service, at 0.facebook.com, for unlimited amounts of time. Furthermore, there is now a Facebook client available that works through a phone's SMS text messaging system, for phones that are not web-enabled, which allows even the cheapest phones without any access to the web to use Facebook.

For activist organizations seeking to reach and mobilize large numbers of people who do not have home broadband access or their own PC, this development has concrete implications. Imagine that a human rights organization in Sudan uses open-source blog-hosting software to set up its own website on a secure service, through which its members can completely control their identities and maintain anonymity, and makes sure that visitors to the site are not tracked. The organization also has a Facebook page, which risks shutdown by Facebook staff if the page's creators fail to comply with Facebook terms of service requiring the use of a real name. Facebook's privacy policies and features can also change without warning, exposing activists who take real risks of reprisal by authoritarian regimes when they use their real names on Facebook. In a country where open web browsing on the mobile phone costs significantly more than Facebook access, that organization is forced to work through Facebook—even if they prefer not to, due to security concerns—because the audience it most needs to reach has been driven by economic incentives away from the open web and onto Facebook.

The lack of neutrality in the mobile space also has troubling surveillance implications, because mobile phones are much more closely tied to the specific Internet user than PCs using broadband. Phone companies, which double as ISPs in many countries, can identify the user's device through its international mobile equipment identity (IMEI) number. The connection is identified by the SIM card number as well as the phone number. Even worse, many countries require registration of national ID or passport with the purchase of mobile phones (including SIM card and IMEI numbers), guaranteeing close association of legally obtained phones with a real person's identity. Thus the mobile service provider can track who is accessing Facebook and when.

Mobile carriers are also experimenting with non-neutrality in the United States. In early 2011 the fifth-largest wireless carrier in the United States, MetroPCS, began offering a $40/month "no long-term commitment" service plan on one of its 4G phones, targeting lower-income customers. The service allowed unlimited talking, texting, web browsing, and YouTube access. Other premium multimedia services were available at additional cost, although some popular brand-name services like Skype and Netflix were excluded altogether. Apps from those services were blocked from being downloaded to the user's handset. Though the logic for the blockage is that these tools tend to use a lot of bandwidth, which costs the providers money, the precedent of blocking certain applications and not others is potentially a slippery slope. Because the MetroPCS service appears to violate the FCC's net neutrality rules, MetroPCS and Verizon joined hands in a lawsuit challenging those rules.

If net neutrality is defeated in the United States, what are the longer-term political implications? In addition to the reduction of choice, we also have a question of whether such policies will drive political discourse and activism to a handful of the biggest social media brands, whose services are likely to be chosen as preferred and free services targeting down-market customers. If I am staging a protest somewhere, and I use my camera phone to capture the scene of my friends being manhandled by police, if I want the video to be seen by the greatest number of people, these types of tiered and preferred services are going to force me to upload that video onto YouTube even if I want to use another video-sharing site whose terms of service I might prefer.

If the protest video I have posted to YouTube is removed by company staffers because it includes police violence and thus violates YouTube's community guidelines prohibiting graphic violence (there are numerous documented cases of activist videos being removed for this reason), or due to my accidental use of footage in which copyrighted music is playing in the background, which then triggers a takedown on copyright grounds (another frequent occurrence on social media sites), I am out of luck, particularly if I am trying to mobilize action and public attention around a time-sensitive matter and I do not have time to find a computer with a fixed Internet connection. I cannot use alternative video-sharing websites because my service allows me to use only YouTube.

Big Brother Apple

Apple's app censorship problem reaches well beyond China into unexpected places. In March 2010 Apple shut down, without notice, an iPad application for *Stern*, one of Germany's biggest magazines. It had published erotic content in the printed magazine, and that content along with all other magazine content was automatically duplicated in the iPad app. This content is perfectly legal in Germany, but because some pages of a specific issue were deemed in violation of Apple's app standards, the entire magazine was censored through the app store. Apple told another German magazine, *Bildt*, that it had to alter content if it wanted to keep its app. That year Apple also censored a cartoon version of James Joyce's *Ulysses*, which contained a few images of nudity, despite the fact that the app was already marked as containing adult content. The app was eventually reinstated, but only after a massive public outcry and widespread media coverage.

Apple also censors controversial political and religious content, including several apps featuring the work of Pulitzer Prize–winning cartoonist

Mark Fiore because they "ridiculed public figures." His app mocking President Obama—containing the type of political humor frequently seen on television and in newspapers—was restored soon after Fiore won his Pulitzer. But apps by other cartoonists were not so fortunate.

Another satirist, Daniel Kurtzman, saw rejection by Apple of two apps meant to accompany his two satirical books, *How to Win a Fight with a Conservative* and *How to Win a Fight with a Liberal*. The apps were programmed to generate ridiculous insults, such as "May a commune of gay, Marxist Muslim illegal immigrants open a drive-through abortion clinic in your church" and "Listen, you bongo-playing vegan, if ignorance is bliss, you must be one happy liberal." After repeated calls and e-mails, an Apple employee eventually explained that this app had been rejected because it involved insults directed at various groups of people. That the insults were completely ridiculous, and were actually meant to demonstrate that stereotype-based slurs are idiotic, was lost on Apple.

In response to widespread controversy over app censorship, in late Apple publicized its app review guidelines and established a review board so that developers would have a more systematic way to appeal decisions made against their apps. But censorship complaints continue, from both the Left and the Right of the political spectrum. In March conservative Christian groups erupted in outrage after Apple, responding to a protest petition signed by 150,000 people organized by Change.org, pulled an app created by the Christian group Exodus International that aimed to help "cure" gay men of their homosexuality. A similar app called the Manhattan Declaration, which condemned homosexuality as "immoral," had also been censored by Apple the previous November, also in response to a Change.org petition drive. The National Religious Broadcasters slammed Apple's decision as "a deliberate act of censorship of the religious speech and thought held by mainstream evangelical Christians in America."

Whatever one thinks about the content and message of the apps in question (I personally find them disturbing), the decision-making process behind their removal is troubling from a political and civil liberties standpoint. Though they can be swayed by public opinion if enough people succeed in galvanizing it, companies are also susceptible to the problem James Madison warned against in *Federalist* No. 10, when he warned that without adequate constitutional safeguards, unpopular speech or dissent could easily be silenced by an "unjust and interested majority." In Appledom, whistle-blowers and purveyors of highly controversial—yet constitutional—political views are thus every bit as vulnerable to takedown as the makers of the Exodus International and Manhattan Declaration apps.

Constitutional norms govern people's free speech rights in public spaces, but not in Appledom. Though these two apps are appalling to anybody who believes in gay rights, and though the public expression of homosexuality is offensive to some conservative Christians, the free speech rights of people with both sets of beliefs are protected by the First Amendment. Organizations promoting either side have a constitutionally protected right to publish websites. Even if no company agrees to host your web content, with sufficient technical know-how, anybody can set up a server in their own office or home and host the content themselves. Defending a home-served site from malicious denial of service attacks is increasingly difficult, but the point is that as long as the open web remains the primary platform for Internet content, it is technically possible for controversial groups to reach an audience—as long as they reside in a country where free speech is protected constitutionally and by the legal system. But if the bulk of the potential audience abandons the open web for closed, app-based systems, then political activists are increasingly hostage to the whims of corporate self-governance.

Fortunately for consumers, there are alternatives to Apple devices. The world's best-selling smart-phone platform as of mid-2011 was Google's open-source Android mobile operating system. Unlike Apple's iOS operating system, which runs only on Apple devices, Android can be adapted to run on any device made by any smart-phone or tablet PC manufacturer. Unlike

Apple, which has a strict vetting system for iPhone and iPad apps, which are available only through its official app stores, Android apps can be created and distributed by anybody with the technical ability to do so, which means that Google does not exercise the same sort of political and cultural censorship that Apple imposes on its apps (although controls are still possible at the level of device-maker and local service provider). Google does carry out security vetting for Android apps for inclusion in its official Android App Market service, but unofficial apps are readily available through other third-party services.

As with many aspects of online and offline society, the flip side to more freedom can be reduced security—which is of course why security is government's most compelling justification for depriving people of freedom, and why finding the right balance between freedom and security is among the greatest challenges for any system of governance. This familiar problem now extends into the new dimension of mobile applications.

In mid-2010 the mobile security firm Lookout discovered an Android app providing customized "wallpaper" images that enable the user to personalize the appearance of their smart-phone display. This capability seemed innocuous enough, except that the app was also collecting the user's personal data and sending it to a "mysterious site in China." Google removed more than fifty apps from its Android Market in March 2011 and thirty additional apps in June. Security experts discovered that some of these apps provided seemingly useful or entertaining services but also contained malware that collected personal information from people's phones. Others were pirated imitations of legitimate programs issued by established companies, but reworked to include malicious code. Apple has, not surprisingly, sought to capitalize on this problem, calling Android apps "inferior." In the end, Google cannot control what a user installs on an Android device, because it does not control the device's hardware. Though Apple is not immune to inevitable security breaches, they are easier to contain because Apple controls both its hardware and its operating system, and the company can quickly push out patches and updates to all iPhones and iPads (as long as the user regularly syncs the device with the "mother ship" via iTunes).

Orwell's *1984*—which Apple once used so cleverly to market its products as rebellious and empowering—is the story of a government that manipulates citizens' fear of war and chaos to justify depriving people of freedom of speech and choice. As it has aged, Apple has become more like Big Brother than anybody could have imagined three decades ago.

Of course, the difference between totalitarianism and democracy hardly amounts to a choice between security and criminality. The distinction between democratic governance and anarchy is that government is empowered by citizens to defend them from criminals and malicious attacks, domestic and foreign. The point of "consent of the governed" is that citizens trade absolute freedom to do anything and everything they please regardless of its impact on others, for a certain degree of security for themselves, their families, and their communities. This trade-off works reliably in the citizens' interest, however, only if government is held accountable—not only by a competitive political system and a strong, independent legal system, but also by constitutional constraints that guarantee respect and protection of citizens' most basic rights, which include free expression and assembly.

In governing our access to and use of applications, Apple provides a valuable service by shielding us from malicious criminals. But it also shows troubling disregard for our political rights as citizens. Like all sovereign powers, Apple and other companies that are failing to protect and respect our rights to free expression and assembly are unlikely to change their attitudes or actions unless and until their customers and users force change upon them—just as citizens have forced change on governments that ultimately realized that their survival required it.

2.4.3 TECHNICAL CODE AND THE SOCIAL CONSTRUCTION OF THE INTERNET

ANDREW J. FLANIGAN, CRAIG FLANIGAN, AND JON FLANIGAN

The social constructivist perspective on technology is discussed in several of the readings in Part 1 of this book (1.1.2 Cowan, 1.2.2 Feenberg, and 1.2.4 Wajcman). These readings emphasize that the design of artifacts is a function of a complex interaction among social, cultural, technical, and economic factors. Andrew Feenberg coined the term *technical code* to describe the underlying values and norms that are reflected in technological design.

In this selection, the authors apply Feenberg's theoretical concept to the analysis of the Internet. The Internet can be understood both through its technical aspects—for instance, its decentralized physical structure and its end-to-end design using TCP/IP (Transmission Control Protocol/Internet Protocol)—and through the kinds of social organization and behavior that it makes possible. The decentralized architecture of the Internet, which arose from its original purpose as an information network for use by the U.S. military, has evolved into a globally interconnected network, hosting all forms of communication and accessible to anyone with an Internet connection. The basic design of the Internet implicitly values interactivity, two-way communication, and distributed, decentralized control over information content. With the advent of Web 2.0, the intrinsic openness of Internet communications has facilitated increasing innovation, enhanced agency, personalization, and social networking.

However, competing values that emphasize social and political control as opposed to open access are also shaping the evolution of the Internet. Laws intended to shield children from pornography and online predators, to protect intellectual property, and to censor politically sensitive information are designed to restrict the informational freedom intrinsic to the Internet's design. The technical code of the Internet also includes features that enable governments and large corporations to exercise control over content and restrict access.

Current debates over Internet freedom, net neutrality, and censorship are manifestations of the deep ambivalence in the present technical code of the Internet. How these debates are resolved will determine whether or not the Internet's potential for enhanced personal empowerment and freedom will be realized.

FOCUS QUESTIONS

1. Discuss two ways in which you feel the Internet has contributed to a more open and democratic society. Provide examples supporting your position.
2. Explain how the relationship between the technical properties and social values inherent in an innovation can be better understood through an analysis of the associated technical code. Discuss the various factors that need to be considered.

Source: Reproduced by permission of SAGE Publications, London, Los Angeles, New Delhi and Singapore, from Andrew J. Flanagin, Craig Flanagin, Jon Flanagin, "Technical code and the social

3. How has the "end-to-end" structure of the Internet and its impact on innovation changed our world? Why do you believe those changes are either beneficial or harmful?
4. Discuss key aspects of the Internet that make it possible for technology to be either empowering or enslaving. Provide examples of how it has functioned in those opposite ways.

KEYWORDS

collaborative filtering, direct democracy, empowerment, "folksonomies," HTML, HTTP, interactivity, interconnectivity, Internet, net neutrality, open source, peer-to-peer technologies, reverse engineering, sense of agency, social construction of technology, social movements, social values, software, TCP/IP, technical code, technical design, Web 2.0, Wikipedia

The *technical code* (Feenberg, 1992, 1995a, 1995b) describes the values or assumptions that become manifest in technological design. According to Feenberg (1995a), economic and social interests constitute 'a background of unexamined cultural assumptions literally designed into the technology itself. I call these assumptions the "technical code" ' (p. 87). A technical code analysis, therefore, reveals the underlying assumptions and choices that become built into technologies, which would otherwise remain largely obscure.

By invoking a technical code perspective, this article undertakes the task of 'reverse engineering' the internet in order to make explicit the assumptions, values, and norms that have to date been built into it. In light of its remarkably rapid development in the last decade or so, coupled with its tremendous social, cultural, and economic significance, it is particularly important to understand the internet's evolution, current form, and future potential. A technical code perspective is valuable in this pursuit (Flanagin et al., 2000) and is best re-applied periodically, particularly to technologies with dramatic impact during periods of brisk change.

THE UTILITY OF THE TECHNICAL CODE PERSPECTIVE AND THE SOCIAL CONSTRUCTION OF TECHNOLOGIES

It is well understood that technical design and social values are interrelated (Abbate, 1999; Bowker et al., 1997; Friedman, 1997; Suchman, 1994). Social constructivist perspectives (Bijker et al., 1987; Bijker and Law, 1992; MacKenzie and Wajcman, 1985), for example, argue that technological design is a function of interconnected social, cultural, technical, and economic factors. Technological artifacts thus result from a complex interaction between technical capabilities and the interests and values of many individuals, groups, and organizations. In this manner, technologies reflect felt social needs and current technical capabilities (MacKenzie and Wajcman, 1985) and are both the result of and impetus for social behaviors.

Given their focus on the complex interactions among diverse interest groups and technical capabilities, such perspectives are useful in making sense of the origin and evolution of a wide range of technological artifacts. Indeed, constructivist studies address such diverse technologies as refrigerators (Cowan, 1985), missile guidance systems (MacKenzie, 1987), the electric light (Hughes, 1985), and the bicycle (Pinch and Bijker, 1987). Social-historical constructivist examinations therefore provide fruitful means by which to explore the competing factors that contribute to form material artifacts.

Nonetheless, the relation between technical properties, values, and social outcomes is often obscure. As Feenberg (1995b) notes, 'design standards are controversial only while they are in flux', suggesting that the form of technologies remains largely unnoticed except during rare times of controversy or redesign (p. 15). A technical code perspective is a powerful analytic tool because it exposes this relation by correlating technical properties with the social values leading to them or facilitated by them. In this manner it extends

social constructivism by exploring the assumptions that are designed into technologies as a way to illuminate the values and choices that become manifest in them. Even if the values themselves change over time, a technical code analysis remains relevant, since values are always interpreted within a specific social and cultural context. Thus, one contribution of the technical code perspective is its utility in exposing the social and design consequences and priorities of technologies that might otherwise go unnoticed.[1]

A second contribution is the technical code perspective's emphasis on change. Because a technical code analysis accentuates the coevolution of technical and social change, it is particularly well-suited to highlight shifts in social values and priorities during periods of rapid technical change, as with the internet's swift evolution. Changing patterns of technical code elements are thus indicative of the changing priorities and values in a group or society, and as such warrant special attention.

Related to this, a third contribution of the technical code perspective is that it draws attention to possibilities that have been foreclosed. By helping to articulate the values manifest in current technologies, attention is simultaneously directed at values or assumptions that are *not* realized, via the technical possibilities that do not come to fruition. In the same way that obsolete technical artifacts and outdated social practices instruct contemporary life, technical foreclosure highlights collective priorities and decisions, past and present. As Bijker and Law (1992) note, the technologies we recognize today 'might have been otherwise' (p. 3), suggesting that current and alternative forms are both consequential in their designs and in what they imply (Postman, 1992; Winner, 1986).

Nonetheless, different, and even highly contradictory, technical forms can coexist in the same time and place. For instance, given sufficient demand and tolerance, Hummers and hybrid cars may not only coexist but both thrive in the same consumer market, in spite of their vastly different environmental and economic implications. Whatever the source or nature of differences, however,

public policy (e.g. manufacturing safety standards for children's toys) or, less formally, endure due to user demands (e.g. recycled and recyclable consumer products).

Change, growth, and atypical events draw special attention to the technical codes of artifacts, prompting re-evaluation of their appropriateness and fit. For example, Ralph Nader's (1965) notorious book *Unsafe at Any Speed* revealed design flaws behind the Chevrolet Corvair and, more widely, an auto industry that marketed and sold cars despite knowing that they would cause completely preventable deaths. The ensuing controversy caused attention to be directed to the auto industry at the level of technical code—prior to the controversy, the technical code included the largely unspoken assumption or norm that a certain casualty level was acceptable with automobile operation. Nader's work, however, can be viewed as a *technical code document* that prompted a re-evaluation of consumer safety and, ultimately, changed the form of the auto industry in the US.

Although the technical code perspective as articulated by Feenberg focuses on material technological artifacts, it can also be extended to consider social arrangements such as the structures, processes, and practices that facilitate human interaction. Doing so explicitly recognizes that technologies are physical objects as well as social practices. Working from this view of technology as both tool (e.g. artifacts such as hardware and software) and technique (e.g. procedures, practices, and social/organizational arrangements) enables a more complete view of the values and assumptions embedded in technologies. Thus, not only can underlying values and norms be extracted from material artifacts but, in the same manner, the social arrangements that constitute technologies (e.g. bureaucracy; see Beniger, 1990) also indicate particular assumptions and values.

In fact, the material and social are often inseparable. For example, the open source software movement reflects as much of a social arrangement as a technical structure: absent strong norms of reciprocity and cooperation, the

sharing could not result in viable software that is freely available, continually improved, and openly shared. Similarly, absent the means of rapid and efficient communication and information sharing, the most altruistic intentions would be unlikely to result in massively co-authored and widely available software programs. Indeed, it is the interaction between the material and social that constitutes complex technologies such as the internet.

THE TECHNICAL CODE OF THE INTERNET

A technical code analysis of the internet[2] necessarily implicates its design features, both through an explicit consideration of values solidified in the internet's structure and through reflection on the implicit values suggested by this structure. Moreover, design decisions support particular social features. In the case of the internet, for example, enhanced individual agency and a renegotiated sense of individual and collective authority are privileged. These social features, in turn, support certain types of outcomes, which would not otherwise be possible. Many of these are novel, or at least unprecedented in scope and form, and are created by the affordances stemming from the internet. In the end, however, such affordances are not unchecked: genuine social, political, and technical limits curb their potential. Overall, a technical code analysis of the internet reveals these and other related possibilities, each of which will be discussed in the following sections.

Technical code by design

Technological design is an important indicator of technical code since design features directly reflect the underlying values and assumptions manifest in a technology. Design features may include the physical form of the artifact (e.g. hardware or physical structure), social and operational procedures that define its use (e.g. software or rules of conduct), and the intersection of these. Broadly speaking, design features encompass the way a technology looks as well as how it works, both technically and socially.

Explicit Values by Design

Design elements of the internet indicate assumptions and values that have been built into it. For example, given its military roots, values of survivability and performance, as opposed to more commercial concerns such as simplicity and consumer appeal, were initially designed into the internet's structure (Abbate, 1999). Perhaps the most important decision in designing the internet was to follow a set of network principles now called the end-to-end arguments or end-to-end structure (Saltzer et al., 1984). Many later decisions, and many of the internet's novel features, follow naturally from this early choice. The defining characteristic of an end-to-end system is that network 'intelligence' (discrimination and processing functions) exists primarily at the periphery of the network, while the network pathways remain neutral, handling all data traffic identically.

The code that implements the internet's end-to-end design at the most basic level is TCP/IP (Transmission Control Protocol/Internet Protocol), which defines the protocol for data transfer. TCP/IP provides a reliable and efficient means of transmitting information over the internet by breaking up data into manageable information 'packets', addressing them correctly, and then reassembling them upon arrival at their destination (for an account of potential internetworking alternatives to TCP/IP see Abbate, 1999; Felczak, 2007). The TCP/IP internet protocol suite is responsible for ensuring that data are delivered via the internet with fidelity and, importantly, that different physical networks share a common method of data transfer that enables their global interconnectivity. Unlike previous communication systems, the internet protocol suite is designed to be independent of the physical conduit, with the result that any network that can carry two-way data can also accommodate internet data traffic.

Other design features also inform the internet's technical code. For instance, HTML (HyperText Markup Language, the predominant programming language for the creation of

web pages) and HTTP (hypertext transfer protocol, the main method used to transfer web-based information over the internet) were integral parts of the web's design on both an applied and a conceptual level. On the applied level, the hypertext structure enables not only the sharing of information but allows the creation of a flexible and evolving framework as opposed to the confines of a fixed system. Philosophically, the original creators of hypertext technology envisioned it as 'removing the confines of linearity … [where] ideas may branch in several directions, and paths through these ideas are followed and created by the reader who also becomes author' (Jackson, 1997: 2). Thus, interactivity was from the start built into the internet through a dynamic system where users can switch roles and redefine themselves and the media they access (Abbate, 1999). As a result, this multiple role relationship allows for a greater degree of two-way control of content than with traditional mass media.

Given the wide scope of the internet and the multitude of key players in its generation-long development process, its interoperability (as opposed to proprietary models) is remarkable. Other technologies are by contrast developed as proprietary and have an economic incentive built into them, which increases the chances that one form will succeed at the expense of another. This stands in marked contrast to the way in which the internet has developed, with standards set industry-wide that are for the most part agreed to and followed by all.

The technical code of interconnectivity that results from these design elements is sometimes an emergent feature of users but is at other times a direct result of what Feenberg (1999) calls 'reflexive design', where designers themselves incorporate values into the technical code self-consciously (see also Srinivasan, 2007). Indeed, formal standards-making bodies such as the W3C (World Wide Web Consortium), an international consortium whose mission is to develop web standards, aim explicitly to ensure 'web interoperability' by promoting compatible web technologies that allow seamless interaction between the hardware and software used to access the web (About the W3C, 2007).

Implicit Values by Extension

Extension of the basic data processing principle of end-to-end design has resulted in the evolution from centralized, data processing-intensive networks to the largely distributed, decentralized structure of the internet (see Abbate, 1999). With transition of information processing over time from batch processes to time-sharing devices to personal computers linked by networks, there has been a steady migration of processing power to computers that are under individuals' personal control. So-called 'peer-to-peer' technologies have leveraged this fact to provide wide-scale distributed information sharing among individuals who manage their own storage/processing devices residing at the periphery of the network. Centralized computers merely coordinate information sharing on these peripheral devices, suggesting that control and power have shifted outward, away from central control, toward individuals.

These design features privilege certain things over others, providing additional information about the internet's technical code. For example, providing information online is simpler than verifying it. Everything from personal websites to social networking sites to product ratings on Amazon constitutes a trail of one's preferences, habits, tastes, and personal information. Among other outcomes, this has resulted in new and complex difficulties in assessing the credibility of web-based information (see Metzger et al., 2003). Indeed, information provision is so simple that it has prompted novel privacy issues as well. Whereas historically the chief problem to overcome in information-sharing was how to publicize information effectively, individuals must now take affirmative steps to maintain privacy and control of their personal information. Chat rooms, bulletin boards, online groups and the like routinely archive and make available transcripts of conversational communication that until recently would have dissolved over time or not been easily locatable. Even information that was once largely geographically localized and therefore functionally unattainable by most people, such as address information residing in telephone books, is now readily available online for anyone to access.

THE SHIFT FROM AUTHORITY TO AUTHORITIES: NEW AGENCY FOR INTERNET USERS

The internet's end-to-end structure has another important characteristic: It is especially well suited to *innovation* on the scale of the individual user. For example, the internet was intentionally designed in a manner that allows those with the requisite skills to program new features directly into it (Abbate, 1999). Because the network is largely indifferent to specific data content, and the information processing functions reside at the periphery of the network in the hands of interconnected users, its design suggests an enormous capacity for personalization, through customized information seeking and specific information provision facilitated by the sheer scale of users online.

Indeed, in its brief history the internet has undergone a dramatic shift in the level and nature of individual involvement. There has been a steady appreciation of the hallmark of digital media: the ability to connect to others and share information easily, quickly, and across distance. Importantly, this appreciation has bloomed in parallel to significant technical developments in both access devices (e.g. computer technologies becoming simultaneously more powerful and less expensive) and infrastructure improvements (e.g. the wide availability of broadband connections). Moreover, extensibility (i.e. software and infrastructure design that anticipates growth) suggests that innovative features and uses have been anticipated from the internet's inception. The end result has been the recognition of the power of the internet to take full advantage of its enormous scale in ways that capitalize on the diversity and knowledge of its users (see, for example, Jenkins, 2006).

Among other things, this shift toward greater individual interconnectivity, personalization, and innovation has upset traditional one-to-many models of mass communication. For example, medium-sized to mass audiences were until quite recently accessible exclusively to those who controlled the substantial centralized media apparatus necessary to reach them. Now, however, large-scale audiences up to the size of social movements as well as highly targeted, specialized audiences have been brought into range for meaningful group participation. Indeed, some argue that we may be witnessing the end of mass communication as traditionally conceived, given the substantial opportunities for targeted communication between individuals and groups of widely diverging sizes, across a variety of topics (Chaffee and Metzger, 2001).

Functionally, the emergence of viable alternatives to centrally controlled media outlets and communication venues, coupled with an appreciation and incorporation of the resources of users, signals a broader shift from an information environment in which there existed few, central, relatively inaccessible, well-known, trusted authorities to multiple, decentralized, widely accessible, emergent, competing voices. Ultimately, this shift from single authority to multiple authorities implies an enhanced perception of individual *agency*, or the capacity for individuals to act or exert power.

INTERNET TOOLS, SOCIAL PRACTICES, AND THE OUTCOMES OF ENHANCED AGENCY

Related to this increased sense of agency are changes in the structure, form, and instances of collective activity that invoke internet tools in substantial ways. An enhanced sense of agency has expanded possibilities for realizing collective goals by undercutting several of the major obstacles to successful cooperative behavior. In fact, theorists have recently proposed that collective action theory itself fails to sufficiently account for current collective efforts being facilitated through the use of internet technologies (Bimber et al., 2005; Flanagin et al., 2006; Lupia and Sin, 2003).

Specifically, theories of collective action traditionally posit that individuals face a distinct decision to either 'free-ride' on the efforts of others' contributions toward some 'public good', or to contribute to the good themselves (Bimber et al., 2005). However, with public goods consisting of web-based information (e.g. databases, information repositories, and archives that

become viable when information is provided by many), individuals may in fact contribute their information unintentionally, unwittingly, or even unknowingly. As mentioned earlier, internet and web technologies actually require users in many cases to take proactive steps to withhold information from public view (including everything from discussions on electronic bulletin boards to feedback offered as part of online rating activities to the identities of friends and common interests in social-networking environments). In such cases, the free-riding construct—premised on conscious and explicit decisions by users as to whether to contribute to a public good or not—is not useful in explaining the initiation and sustenance of many forms of contemporary collective behavior.

In addition, formal organizations have typically been viewed as necessary for coordinating people's collective action efforts by motivating them to make private resources publicly available and persuading them to remain involved in collective efforts despite substantial risks and potential setbacks (Bimber et al., 2005). Indeed, in many instances, formal organizations are indispensable for overcoming communicative and administrative challenges in the orchestration of efficient collective effort. Yet, instances of coordinated collective action without formal organization are increasingly common when efforts rely in various ways on internet technologies. From political advocacy efforts such as the US presidential campaign of Howard Dean in 2003 and the political action organization MoveOn.org, to self-organizing networks of 'smart mobs' that mobilize people to gather at a particular moment for a common goal on short notice (Rheingold, 2003), to the open source software development movement, the internet has enabled a diversity of collective efforts that do not fit well into traditional collective action theory.

Thus, collective action theories as traditionally formulated are most useful under conditions where information sharing boundaries are firm and comparatively impermeable, and individuals' efforts to cross them are characterized by discrete free-riding calculations in the context of high costs (Bimber et al., 2005; Flanagin et al., 2006). Formal organizations are required for collective action in circumstances where costs of information, communication, and coordination are high. Although such conditions remain important, they are no longer universal given the substantial affordances provided by the internet. When boundaries between private and public domains are permeable and easily crossed, people's negotiation of the boundary typically involves less intentionality and calculation and the formal structures designed to broker the private-to-public transition become less crucial. Indeed, such a transition may even be inadvertent, or it may involve a more fluid movement back and forth between public and private domains of interest and action. This further emphasizes individuals' reinvigorated sense of agency, given the expanded possibilities they now have for realizing their collective goals, and suggests that the technical code of the internet exhibits an enhanced sense of *empowerment*, inasmuch as internet users feel their collective efforts result in desired outcomes.

The notion of enhanced agency is also evident in current rhetoric about so-called 'web 2.0' functionality and philosophy, which is itself highly indicative of the internet's current technical code. Although there is some disagreement on its definition (O'Reilly, 2005), web 2.0 typically denotes an internet environment in which tools are used to leverage the potential contributions of a wide variety of users, each of whom is believed to contribute value in some manner to collective endeavors. Rather than seen primarily as a static information delivery platform, this view conceives of the internet as a dynamic collaborative environment in which diverse opinions, experiences, and skills can be aggregated to provide substantial resources. The essential premise is that given efficient means of information sharing and participation, collective benefits will emerge from individual contributions.

Examples of this type of wide-scale collaboration abound: For instance, 'credentialing' activity relies on peer rating of some dimension of an individual, organization, transaction, opinion, or product that is subsequently tallied to provide an omnibus rating score. A simple example is the online auction site eBay.com, which relies on its members to rate others with whom they have

engaged in a transaction, in order to mitigate the considerable risk involved in financial transactions with unknown others. Credentialing thus resolves the problem of the lack of reliable, firsthand knowledge by enhancing the distribution of existing information. Other tools similarly recognize and exploit the aggregated experiences of individuals. In order to predict a range of phenomena—from movies, books, or news stories a person is likely to enjoy or find relevant to the degree to which an email message is likely to be spam—'collaborative filtering' invokes a similarity metric, calculated from the preferences of a subgroup of like others, that is used to make recommendations to a target individual. Collaborative filtering thus serves to automate 'word-of-mouth' recommendations, tailored to a person based on his or her individual characteristics by simultaneously radically expanding and making manageable the potential pool of known others.

Further exemplifying the emergence of individual experiences and preferences are tools that promote user self-provision and self-organization of web content. For instance, the English-language version of the online encyclopedia Wikipedia currently boasts nearly 3 million user-created articles that have been judged to be nearly as accurate as well-regarded print encyclopedias, where the content is supplied solely by established experts (Giles, 2005). This collaborative resource suggests the enormous knowledge assets that reside in collectives, which until recently remained largely untapped due to insurmountable coordination costs.

What is called the 'self-organization' of web content has also burgeoned in recent years. For instance, recent online tools have begun to support a shift from taxonomies of web information, where information is organized according to a centrally-imposed relational or organizational structure, to more emergent classifications based on users' own categorization (so-called 'folksonomies', indicating their emergence from individual users). The use of 'tags', for example, where individuals select key words of their own choosing that are in turn used to organize content ranging from websites to more specific pieces of information (e.g. photos, news stories, or ideas) are increasingly heralded as a means of bottom-up organization of web content. Tags can be presented in a way that provides a visually intuitive means of organizing web content (i.e. as 'tag clouds'), leading to hyperlinked associations among items based on the popularity of associations. Importantly, collaborative categorizations based on tags can represent multiple associations that may overlap, rather than mutually exclusive identifiers. Tags also appear to provide a relatively low cost means of user input, thereby overcoming the potential obstacle to collaborative filtering tools, which opens them up to a wide range of applications, from photograph sorting and sharing (e.g. http://flickr.com/) to 'social bookmarking' of a wide variety of web content (e.g. http://del.icio.us/).

Collectively, the reformulation of traditional collective action theory and the various means of individual participation and collaboration suggest the power of the 'long tail' distribution of internet user preferences. Due to reductions in the costs of production and distribution stemming from contemporary technologies, Anderson (2006) argues that substantial opportunities exist to tap into that portion of the demand curve represented by products that may be low in popularity overall, but that collectively represent a large proportion of the market (i.e. those that reside in the 'long tail' of the distribution curve). This idea of the long tail extends beyond consumer product demand. As evidenced by new forms of formalized collective action efforts, as well as the less formal collective behaviors of credentialing, collaborative filtering, and the self-provision and self-categorization of information through tools such as wikis and tags, the potential of the long tail of internet users is enormous. As communication and information sharing costs are reduced, new opportunities that take advantage of the resources of collectivities—in the form of everything from the provision of consumer information to political protests—arise (Benkler, 2006). Moreover, it appears that the current generation of information sharing tools mark only the early stages of effectively tapping into the specific skills and knowledge of the internet's existing user pool.

From a technical code standpoint, the emergence of collective action enhanced by internet connectivity and so-called web 2.0 affordances, behaviors, and philosophy signal increased agency, as noted earlier, as well as the values of democratic participation, empowerment, collaboration, and social engagement. Indeed, the acknowledgment (and labeling) of web 2.0 itself signals a wider recognition of the internet's evolving technical code, inasmuch as the articulation of web 2.0 characteristics and philosophies indicates an underlying recognition and appreciation of the values and assumptions that are increasingly manifest in the internet's current usage and potential.

LIMITS TO INTERNET OPENNESS

Elements of the internet's technical code articulated so far—decentralized technical interoperability that increases the potential for innovation, coupled with a sense of empowerment achieved through enhanced agency and collaboration—imply that the internet is resistant to centralized control and is not adversely affected by impediments to these outcomes. Nonetheless, although these technical code features legitimately emerge from an analysis of the internet in its current form, there also exist countervailing forces, which are also indicative of the internet's technical code. Indeed, as noted earlier, technical codes can co-exist in oppositional ways, and the very technical features that currently appear to engender relative freedom can also be employed to exert strict control. In the end, conclusions about the internet's technical code must account for both forces.

Many internet-specific laws in the US illustrate this opposition. Given the open access architecture and inclusive reach of the internet, several laws have been enacted specifically to inhibit the internet's openness, particularly in cases where this openness implicates protected audiences such as children. For instance, the Communications Decency Act (CDA, 1996) and the Child Online Protection Act (COPA, 1998) were enacted in order to shield children from pornographic material on the web. Similarly, the Child Pornography Prevention Act (CPPA, 1996) attempted to expand the definition of child pornography to include images that appear to depict children but actually do not, including images of youthful-looking adults or images that are computer-generated. The Deleting Online Predators Act (DOPA, 2006) would require schools and libraries to restrict minors' access to online 'commercial networking sites' and chat rooms, in order to retain subsidized funding for their internet access. Finally, the Children's Internet Protection Act (CIPA, 2000) requires libraries and schools to install filters on their internet computers designed to block obscene or pornographic material, or material deemed 'harmful to minors'. These laws are complex and have been the target of considerable legal challenge, with results ranging from being largely overruled (e.g. CDA, CPPA, and COPA), to being upheld (e.g. CIPA), to battles that are ongoing (e.g. DOPA). Nonetheless, these acts and laws collectively illustrate specific efforts to counteract the internet's information sharing capabilities, at least with regard to a specific 'vulnerable' audience. Notably, each serves as a technical code document, inasmuch as they all articulate specific social values.

Legal efforts to thwart the open access inherent in the internet's design extend beyond protected audiences like children. For example, as noted earlier, one outcome of the internet's end-to-end design is the potential of peer-to-peer communication and information sharing. The heart of peer-to-peer architecture is the ability to efficiently connect end-users directly with one another, in order that people might share information among members of widely distributed groups. In the late 1990s peer-to-peer file sharing became popular among individuals with access to the internet. Perhaps the most highly-publicized instance of wide scale file sharing at the time was the peer-to-peer service Napster, which facilitated individuals' exchange of MP3 format music files, among other things. In 1999 Napster was charged with contributing to copyright infringement by its users, and was forced to cease operations as a free file-sharing peer-to-peer system. More recently, in the case of MGM versus the peer-to-peer software company Grokster (545 U.S. 913), the US Supreme Court ruled that internet file-sharing

services will be held responsible if they *intend* for their customers to use software primarily to swap songs and movies illegally. In addition, the Digital Millennium Copyright Act (DMCA) makes it illegal to circumvent technological copyright protection mechanisms that can be used on CDs or DVDs or to distribute tools that enable such circumvention. In these ways, the widespread and simplified exchange of copyrighted information, at least as it once existed in the US and elsewhere during Napster's heyday, is one potential technical and social feature or value that has been legally—if not functionally—foreclosed.

In essence, the conflict between peer-to-peer file sharing services and the recording industry represents a conflict at the level of the technical code: although peer-to-peer file sharing was, and remains, a key capability (value) of the internet, it is at odds with the goal of some copyright holders to enforce their exclusive rights to intellectual property (a conflicting value). Specific enforcement of copyright law in this instance, and in the manner executed to date, can thus be seen as a countervailing force designed to reign in the inherent openness of the architecture and design of the internet. For example, the new 'inducement' doctrine of copyright infringement liability arising from the *MGM v. Grokster* case can only serve to inhibit innovativeness. Ultimately, this conflict draws attention to the internet's technical code by highlighting the tension surrounding the innovativeness designed into the internet and which values and assumptions are being realized. Interestingly, efforts to resolve or mitigate this conflict have arisen through the internet itself. For example, 'creative commons' licensing, designed to enable copyright holders to stipulate and publicize specific legal uses of their intellectual property, relies on the internet's information sharing capacity to enact a system of more flexible copyright than has existed before.

These conflicts begin to illustrate the complexity of assessing technical codes, particularly when artifacts are intricate, various interests are involved, and boundaries between groups are fluid or widely inclusive. Although well-matched within particular social environments, technical codes that clash with widely-held principles, ideologies, or beliefs are likely to result in conflict, as is often the case when technologies cross group boundaries. Indeed, this type of conflict may be more pronounced with some technologies than others. Given the internet's global reach, it seems more likely than many technologies to prompt meaningful and instructive technical code conflicts.

For example, as a global technology, the dominant technical code of the internet is bound to violate political, ideological, and cultural norms and expectations across countries. One prominent instance is the attempt by the Chinese government to control the information accessed by its citizens via the internet. In essence, the internet's technical code conflicts with priorities and controls that the Chinese government wishes to impose. Thus, the Chinese government is faced with the task of re-aligning the internet to conform to its view of how communication and information sharing tools should be used. However, given the internet's inherent ability to circumvent information loss or single points of control, close monitoring of the 135 million people currently using the internet in China is a daunting task (Fowler, 2007). Nonetheless, the Chinese government has employed sophisticated filtering mechanisms that prohibit all but the most ambitious users in China from accessing a wide range of content, including pornographic, religious, and political information (Internet Filtering in China, 2005).

However, the internet address and domain name registration system can greatly facilitate this type of centralized control. Unlike the distributed end-to-end architecture of the internet generally, the technology and administrative control of internet and Web addresses and names is highly centralized and hierarchical. This address or naming system, in fact, ties internet activities to particular jurisdictions, agencies, and even individuals. It therefore presents significant potential for centralized control over its more distributed and heterogeneous aspects. For example, nations like Germany can block content promoting Nazi ideology or denying the holocaust within its own borders—even though the

servers providing this content might be located in other nations and jurisdictions. Similarly, this naming system allows the Chinese government to block international news content from the BBC's website, while allowing content from CNN, for example. This system also makes it possible for internet service providers to contravene internet neutrality principles by enabling them to identify the source of a given packet of information, and to expedite or delay its delivery in subsequent switching and transmission procedures. This part of the internet's architecture thus illustrates the ambivalence of values articulated in its technical code, and in the values and priorities that can be articulated through it.

There thus exist several forces working against the internet's decentralized, innovative, collaborative, and empowering structures. Many of these are represented by laws, which constitute a relatively explicit statement of technical code in their expression of the decisions and boundary conditions collectively imposed on a technology. Other countervailing forces are social, such as political and social pressures to enact laws in the first place, and normative uses that dictate acceptable and unacceptable behaviors online. As the economic, legal, and social influences on the internet converge, innovative solutions that balance freedom with control are also emerging (e.g., creative commons copyright licensing), many of which in fact take fundamental advantage of the internet's distinctive features.

THE ROLE OF (TECHNICAL) CODE IN REALIZING THE INTERNET'S POTENTIAL

Overall, a technical code analysis of the internet reveals a decentralized interoperable architecture that is constrained by certain countervailing forces. Technical characteristics that enhance individuals' agency, and increase the likelihood of collaboration and innovation, also present the potential for centralized control. As Castells (2001) argues, the internet represents a complex 'double game,' where the very tools used to promote democratic participation can also be used to restrict it. This ambivalence in the internet's technical code illustrates both its enormous potential to realize innovation and freedom as well as its facility for constraining these.

There is evidence that the positive potential promised by the internet is being realized. Successful political advocacy efforts are enabled by online organization and mobilization (Earl, 2006); bottom-up, participatory culture relying in large part on the internet can increase collaboration and enhance the communal practice of cultural production and consumption (Jenkins, 2006); and group 'intelligence,' collated and tallied with the aid of web-based tools, can in many cases lead to superior solutions to a wide range of problems (Surowiecki, 2004). Moreover, many of these efforts result from a tacit awareness of technical code principles: instances of reflexive design offer hope that communities can create information systems tailored to their own particular cultural contexts in order to satisfy their own priorities (Srinivasan, 2007). In the extreme, even 'direct democracy' has been proposed, where agency and empowerment coincide through citizens' direct expression in law, trials, and decrees or motions by all those who wish to participate in their own governance. Such efforts take reflexive design one step further with a transparent convergence of design and values, which is rarely the case.

The current technical code of the internet suggests the partial realization and continued potential for these various types of relatively unrestrained participation and collaboration. Indeed, it is precisely this potential the Chinese government is vigorously attempting to restrain among its citizens. As Pool (1983) observed relatively early in the development of networked technologies, the potential of electronic media is monumental:

> They allow for more knowledge, easier access, and freer speech than were ever enjoyed before. They fit the practices of print. The characteristics of media shape what is done with them, so one might anticipate that these technologies of freedom will overwhelm all attempts to control them. (p. 251)

Yet this potential for freedom and the internet's ability to resist control are more fragile than the overall evidence from a technical code analysis

might at first suggest. Structural features facilitating individual agency and freedom also present a 'relatively easy target for political and commercial exploitation' (Beniger, 1996: 59; see also Castells, 2001). As Lessig (1999, 2006) notes, the 'code' of the internet, or the software and hardware that regulate and thus govern cyberspace, is malleable and therefore open to control by various forces. In fact, the same technical features supporting open access can be applied to exert various forms of control. For example, digital networked technologies provide remarkably efficient and accurate surveillance mechanisms that can seriously violate individuals' privacy, particularly if coupled with reliable digital identification, which some argue is on the near horizon (Lessig, 2006).

Thus, the same tools that enable increased individual agency can also be used to diminish it through deceptive and carefully crafted ploys that exploit individuals' security, privacy, or freedom of expression. Managing the balance between these forces will be one of the critical challenges over the next several years and, depending on what forces are privileged, future technical code analyses of the internet might suggest radically different values and assumptions than are now evident. To assess this, periodic reexamination of the technical code is required.

To frame the internet's possible futures, Lessig (2006) describes a tension between so-called 'east coast' code (suggesting the Washington, DC origin of much US law) embedded in laws that dictate what can and cannot be done with the internet, and the 'west coast' computer code (suggesting the Silicon Valley/Microsoft influence), or 'the instructions imbedded in the software and hardware that make cyberspace work' (p. 72). Lessig argues that the power of east over west is increasing: Code is increasingly the product of commercial entities (as opposed to nonprofits or individuals), which by their nature are subject to many levels of control. As some argue, cyberspace will ultimately be controlled by those controlling society at large, namely large corporations with interests that are often at odds with those ascribed, perhaps idealistically, to the early denizens of the internet (Beniger, 1996).

Lessig therefore advocates actively using laws to control the basic architecture of the internet because, left unchecked, the internet is likely to experience 'the opposite of its architecture at its birth. This invisible hand, pushed by government and by commerce, is constructing an architecture that will perfect control and make highly efficient regulation possible' (Lessig, 2006: 4). According to him, 'code' is thus crucial to consider under any future scenario because it will present the greatest threat to both liberal and libertarian ideals, as well as their greatest promise. We can build, or architect, or *code* cyberspace to protect values that we believe are fundamental. Or we can build, or architect, or code cyberspace to allow those values to disappear. (Lessig, 2006, p. 6, emphasis in original)

CONCLUSION

The technical code perspective provides a penetrating view of technologies by increasing awareness of the choices that form them, and by suggesting the consequences of those choices. When attention is drawn to the technical code, public interest and involvement in related policies and practices tend to be lively and robust. For instance, recent debate and action surrounding battles over 'net neutrality,' or the continuation of the internet's nondiscriminatory approach to data traffic, has sparked intense citizen involvement, governmental action, and wide scale media attention. In effect, this debate has taken place at the level of the technical code: The question of net neutrality has fundamentally been framed as 'what values should be built into the internet?' In spite of its importance, attention is seldom focused at the level of the technical code. People rarely look beyond current technological forms to consider alternative versions, past or future, and public attention is rarely directed at the underlying values that are implicit in technical choices. Yet, many suggest that this is precisely the level of public involvement and discourse that is required, particularly during times of rapid technological change. To ensure that the internet accurately reflects individual users' values and preferences, 'there are choices to be made about how this network evolves. These choices will affect fundamentally what values are built into the

network.' (Lessig, 2006: 311). Put another way, what is required is directed and conscious influence over technical design and decision-making, rather than passive acceptance of current (or continuing) forms and tacit endorsement of the forces increasingly serving to control them. With sufficient initiative and participation, users of technologies have at least the potential to *choose* the values that come to reside in technologies through active engagement in their use and (re)design (Feenberg, 1995b). Such choices are not trivial. In the end, they are likely to determine the degree to which the potential of the internet as a 'technology of freedom' will continue to be realized, or will be foreclosed.

NOTES

1. It is important to note that design decisions originate from technology users at least as much as they do from the original 'designers' of technology (e.g., engineers, programmers, and the like). Technologies are constantly modified, reinvented, and appropriated in their use (Abbate, 1999; McPherson, 2008; Jenkins, 2006; Poole and DeSanctis, 1990), often resulting in novel forms and unintended consequences. In the process of user appropriation, just as in the process of initial design execution, values become embedded in technologies. A technical code analysis can help to reveal which among these values have been privileged, at least for the time being.
2. Although the 'internet' refers to the physical infrastructure of interconnected computers, cables, and other devices that serves as the infrastructure for global communication, it is often conflated with a host of other services communicated via the internet, such as the web. From a technical code perspective, 'the internet' is best conceived of as the internet (i.e. an infrastructure) and its associated applications (e.g. the web), as well as a series of design and user decisions. This more inclusive definition best highlights the internet's sociotechnical nature.

REFERENCES

Abbate, J. (1999) *Inventing the Internet*. Cambridge, MA: MIT Press.

About the W3C (2007) URL (consulted 3 January 2007): http://www.w3.org/Consortium/

Anderson, C. (2006) *The Long Tail: Why the Future of Business Is Selling Less of More*. New York: Hyperion. Beniger, J.R. (1990) 'Conceptualizing Information Technology as Organization, and Vice Versa', in J. Fulk and C. Steinfield (eds) *Organizations and Communication Technology*, pp. 29–45. Newbury Park, CA: Sage.

Beniger, J.R. (1996) 'Who Shall Control Cyberspace?', in L. Strate, R. Jacobson and S.B. Gibson (eds) *Communication and Cyberspace: Social Interaction in an Electronic Environment*, pp. 49–58. Cresskill, NJ: Hampton Press.

Benkler, Y. (2006) *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. New Haven, CT: Yale University Press.

Bijker, W. and J. Law (eds) (1992) *Shaping Technology/Building Society*. Cambridge, MA: MIT Press.

Bijker, T., T. Hughes and T. Pinch (eds) (1987) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. Cambridge, MA: MIT Press.

Bimber, B., A.J. Flanagin and C. Stohl (2005) 'Reconceptualizing Collective Action in the Contemporary Media Environment', *Communication Theory* 15: 365–88.

Bowker, G.C., W. Turner, S.L. Star, W. Turner and L. Gasser (1997) *Social Science, Technical Systems, and Cooperative Work: Beyond the Great Divide*. Mahwah, NJ: Erlbaum.

Castells, M. (2001) *The Internet Galaxy: Reflections on the Internet, Business, and Society*. New York: Oxford University Press.

Chaffee, S.H. and M.J. Metzger (2001) 'The End of Mass Communication?', *Mass Communication and Society* 4(4): 365–79.

Cowan, R.S. (1985) 'How the Refrigerator Got Its Hum', in D. MacKenzie and J. Wajcman (eds) *The Social Shaping of Technology*, pp. 202–18. Philadelphia, PA: Open University Press.

Earl, J. (2006) 'Pursuing Social Change Online: The Use of Four Protest Tactics on the Internet', *Social Science Computer Review* 24(3): 362–77.

Feenberg, A. (1992) 'From Information to Communication: The French Experience with Videotext', in M. Lea (ed.) *Contexts of Computer-mediated Communication*, pp. 168–87. London: Harvester Wheatsheaf.

Feenberg, A. (1995a) *Alternative Modernity: The Technical Turn in Philosophy and Social Theory*. Berkeley, CA: University of California Press.

Feenberg, A. (1995b) 'Subversive Rationalization: Technology, Power, and Democracy', in A. Feenberg and A. Hannay (eds) *Technology and the Politics of Knowledge*, pp. 3–22. Bloomington, IN: Indiana University Press.

Feenberg, A. (1999) *Questioning Technology*. New York: Routledge.

Feleczak, M. (2007) 'The Ghosts of Internet Past? X.25, IPv6, and Control in Cyberspace', in A. Feenberg and N. Friesen (eds) *(Re)Inventing the internet: Critical case studies*. Publisher?

Flanagin, A.J., W.J.M. Farinola and M.J. Metzger (2000) 'The Technical Code of the Internet/World Wide Web', *Critical Studies in Media Communication* 17: 409–28.

Flanagin, A.J., C. Stohl and B. Bimber (2006) 'Modeling the Structure of Collective Action', *Communication Monographs* 73: 29–54.

Fowler, G.A. (2007) 'It's Called the Forbidden City for a Reason', *Wall Street Journal*, 19 January, B1.

Friedman, B. (ed.) (1997) *Human Values and the Design of Computer Technology*. Cambridge: Cambridge University Press.

Giles, J. (2005) 'Special Report: Internet Encyclopaedias Go Head to Head', *Nature* 438: 900–1.

Hughes, T.P. (1985) 'Edison and Electric light', in D. MacKenzie and J. Wajcman (eds) *The Social Shaping of Technology: How the Refrigerator Got Its Hum*, pp. 39–52. Philadelphia, PA: Open University Press.

Internet Filtering in China in 2004–2005: A Country Study. URL (consulted 8 January 2007): http://www.opennetinitiative.net/studies/china/ONI_China_Country_Study.pdf

Jackson, M. (1997) 'Assessing the Structure of Communication on the World Wide Web', *Journal of Computer-mediated Communication* 3(1). URL (consulted 28 February 2007): http://jcmc.indiana.edu/vol3/issue1/jackson.html

Jenkins, H. (2006) *Convergence Culture: Where Old and New Media Collide*. New York: New York University Press.

Lessig, L. (1999) *Code and Other Laws of Cyberspace*. New York: Basic Books.

Lessig, L. (2006) *Code: Version 2.0*. New York: Basic Books.

Lupia, A. and G. Sin (2003) 'Which Public Goods Are Endangered?: How Evolving Communication Technologies Affect the Logic of Collective Action', *Public Choice* 117: 315–31.

MacKenzie, D. (1987) 'Missile Accuracy: A Case Study in the Social Processes of Technological Change', in T. Bijker, T. Hughes, and T. Pinch (eds) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*, pp. 195–222. Cambridge, MA: MIT Press.

MacKenzie, D. and J. Wajcman (1985) *The Social Shaping of Technology: How the Refrigerator Got Its Hum*. Philadelphia, PA: Open University Press.

McPherson, T. (ed.) (2008) *Digital Youth, Innovation, and the Unexpected*. Cambridge, MA: MIT Press.

Metzger, M.J., A.J. Flanagin, K. Eyal, D.R. Lemus and R. McCann (2003) 'Credibility in the 21st Century: Integrating Perspectives on Source, Message, and Media Credibility in the Contemporary Media Environment', in P. Kalbfleisch (ed.) *Communication Yearbook* 27, pp. 293–335. Mahwah, NJ: Lawrence Erlbaum.

Nader, R. (1965) *Unsafe at Any Speed*. New York: Grossman.

O'Reilly, T. (2005) 'What Is Web 2.0?: Design Patterns and Business Models for the Next Generation of Software', URL (consulted 22 January 2007): http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html

Pinch, T.J. and W.E. Bijker (1987) 'The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other', in T. Bijker, T. Hughes and T. Pinch (eds) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*, pp. 17–50. Cambridge, MA: MIT Press.

Pool, I. (1983) *Technologies of Freedom*. Cambridge, MA: The Belknap Press of Harvard University Press.

Poole, M.S. and G. DeSanctis (1990) 'Understanding the Use of Group Decision Support Systems: The Theory of Adaptive Structuration', in J. Fulk and C.W. Steinfield (eds) *Organizations and Communication Technology*, pp. 173–93. Newbury Park, CA: Sage.

Postman, N. (1992) *Technopoly: The Surrender of Culture to Technology*. New York: Vintage Books.

Rheingold, H. (2003) *Smart Mobs*. New York: Basic Books.

Saltzer, J.H., D.P. Reed and D.D. Clark (1984) 'End-to-end Arguments in System Design', *ACM Transactions on Computer Systems* 2: 277–88.

Srinivasan, R. (2007) 'Ethnomethodological Architectures: Information Systems Driven by Cultural and Community Visions', *Journal of the American Society for Information Science and Technology* 58: 723–33.

Suchman, L. (1994) 'Do Categories Have Politics? The Language/Action Perspective Reconsidered', *Computer Supported Cooperative Work* 2: 177–90.

Surowiecki, J. (2004) *The Wisdom of Crowds*. New York: Random House.

Winner, L. (1986) *The Whale and the Reactor: A Search for Limits in an Age of High Technology*. Chicago: University of Chicago Press.

2.5 BIOTECHNOLOGY

2.5.1 OUR POSTHUMAN FUTURE

FRANCIS FUKUYAMA

In this reading, drawn from his book, *Our Posthuman Future: Consequences of the Biotechnology Revolution*, Francis Fukuyama argues that biotechnology may propel us into a future in which some of our most basic political values will be thrown into question. In particular, he raises the specter of a future society in which advances in medicine and biotechnology enable the wealthy to endow their children with socially desirable characteristics that will give them a competitive advantage over other people. Unlike the kinds of advantages that the children of the wealthy currently enjoy—for instance, better educational opportunities—these competitive advantages will be "bred into the bone" and passed on from one generation to the next through the use of new reproductive technologies. We may arrive at a point sometime in the not-too-distant future in which it will no longer be possible to claim that "all persons are created equal" because we will know for a fact that some are not—they are genetically engineered to be superior.

This selection, and the others in this section, raises some of the most troubling and far-reaching moral questions posed by contemporary technology. Should we permit biotechnology to develop so as to allow us to control the design of future human beings? Or should we regulate or possibly renounce these technological powers in order to preserve human values?

Source: "A Tale of Two Dystopias" from *our posthuman future: consequences of the biotechnology revolution* by Francis Fukuyama. Copyright © 2002 by Francis Fukuyama. Reprinted by permission of Farrar, Straus and Giroux, LLC and International Creative Mangement, Partners. All rights reserved.

FOCUS QUESTIONS

1. The author presents three scenarios of how the biotechnology revolution may alter our conception of human nature over the coming years. Which, if any, of these scenarios do you find morally troubling? Why?
2. What are some forms of technology that we regulate intensively? What forms of technology do we allow to develop without much government oversight or regulation? What are the features of technologies that justify these differing regulatory approaches?
3. What conclusion does the author draw from the failure of socialist revolutions in places like Russia, China, Cuba, and Cambodia? Do you agree that with his conclusion?
4. Why does Fukuyama believe that biotechnology has the potential to undermine the idea of human equality?
5. Do the threats that Fukuyama is worried about depend on having "state-sponsored" eugenics programs, which coerce people's reproductive choices, as in the novel *Brave New World*? Explain why or why not.

KEYWORDS

biotechnology, *Brave New World*, class war, cultural plasticity, government regulation, human nature, liberal democracy, posthuman

I was born in 1952, right in the middle of the American baby boom. For any person growing up as I did in the middle decades of the twentieth century, the future and its terrifying possibilities were defined by two books, George Orwell's *1984* (first published in 1949) and Aldous Huxley's *Brave New World* (published in 1932).

The two books were far more prescient than anyone realized at the time, because they were centered on two different technologies that would in fact emerge and shape the world over the next two generations. The novel *1984* was about what we now call information technology: central to the success of the vast, totalitarian empire that had been set up over Oceania was a device called the telescreen, a wall-sized flat-panel display that could simultaneously send and receive images from each individual household to a hovering Big Brother. The telescreen was what permitted the vast centralization of social life under the Ministry of Truth and the Ministry of Love, for it allowed the government to banish privacy by monitoring every word and deed over a massive network of wires.

Brave New World, by contrast, was about the other big technological revolution about to take place, that of biotechnology. Bokanovskification, the hatching of people not in wombs but, as we now say, in vitro; the drug soma, which gave people instant happiness; the Feelies, in which sensation was simulated by implanted electrodes; and the modification of behavior through constant subliminal repetition and, when that didn't work, through the administration of various artificial hormones were what gave this book its particularly creepy ambiance.

With at least a half century separating us from the publication of these books, we can see that while the technological predictions they made were startlingly accurate, the political predictions of the first book, *1984*, were entirely wrong. The year 1984 came and went, with the United States still locked in a Cold War struggle with the Soviet Union. That year saw the introduction of a new model of the IBM personal computer and the beginning of what became the PC revolution. As Peter Huber has argued, the personal computer, linked to the Internet, was in fact the realization of Orwell's telescreen. But instead of becoming an instrument of centralization and tyranny, it led to just the opposite: the democratization of access to information and the decentralization of politics. Instead of Big Brother watching everyone, people could use the PC and Internet to watch Big Brother, as governments everywhere were driven to publish more information on their own activities.

Just five years after 1984, in a series of dramatic events that would earlier have seemed like political science fiction, the Soviet Union and its empire collapsed, and the totalitarian threat that Orwell had so vividly evoked vanished. People were again quick to point out that these two events—the collapse of totalitarian empires and the emergence of the personal computer, as well as other forms of inexpensive information technology, from TVs and radios to faxes and e-mail—were not unrelated. Totalitarian rule depended on a regime's ability to maintain a monopoly over information, and once modern information technology made that impossible, the regime's power was undermined.

The political prescience of the other great dystopia, *Brave New World*, remains to be seen. Many of the technologies that Huxley envisioned, like in vitro fertilization, surrogate motherhood, psychotropic drugs, and genetic engineering for the manufacture of children, are already here or just over the horizon. But this revolution has only just begun; the daily avalanche of announcements of new breakthroughs in biomedical technology and achievements such as the completion of the Human Genome Project in the year 2000 portend much more serious changes to come.

Of the nightmares evoked by these two books, *Brave New World*'s always struck me as more subtle and more challenging. It is easy to see what's wrong with the world of 1984: the protagonist, Winston Smith, is known to hate rats above all things, so Big Brother devises a cage in which rats can bite at Smith's face in order to get him to betray his lover. This is the world of classical tyranny, technologically empowered but not so different from what we have tragically seen and known in human history.

In *Brave New World*, by contrast, the evil is not so obvious because no one is hurt; indeed, this is a world in which everyone gets what they want. As one of the characters notes, "The Controllers realized that force was no good," and that people would have to be seduced rather than compelled to live in an orderly society. In this world, disease and social conflict have been abolished, there is no depression, madness, loneliness, or emotional distress, sex is good and readily available. There is even a government ministry to ensure that the length of time between the appearance of a desire and its satisfaction is kept to a minimum. No one takes religion seriously any longer, no one is introspective or has unrequited longings, the biological family has been abolished, no one reads Shakespeare. But no one (save John the Savage, the book's protagonist) misses these things, either, since they are happy and healthy.

Since the novel's publication, there have probably been several million high school essays written in answer to the question, "What's wrong with this picture?" The answer given (on papers that get A's, at any rate) usually runs something like this: the people in *Brave New World* may be healthy and happy, but they have ceased to be *human beings*. They no longer struggle, aspire, love, feel pain, make difficult moral choices, have families, or do any of the things that we traditionally associate with being human. They no longer have the characteristics that give us human dignity. Indeed, there is no such thing as the human race any longer, since they have been bred by the Controllers into separate castes of Alphas, Betas, Epsilons, and Gammas who are as distant from each other as humans are from animals. Their world has become unnatural in the most profound sense imaginable, because *human nature* has been altered. In the words of bioethicist Leon Kass, "Unlike the man reduced by disease or slavery, the people dehumanized à la *Brave New World* are not miserable, don't know that they are dehumanized, and, what is worse, would not care if they knew. They are, indeed, happy slaves with a slavish happiness."

But while this kind of answer is usually adequate to satisfy the typical high school English teacher, it does not (as Kass goes on to note) probe nearly deeply enough. For one can then ask, What is so important about being a human being in the traditional way that Huxley defines it? After all, what the human race is today is the product of an evolutionary process that has been going on for millions of years, one that with any luck will continue well into the future. There are no fixed human characteristics, except for a general capability to choose what we want to be,

to modify ourselves in accordance with our desires. So who is to tell us that being human and having dignity means sticking with a set of emotional responses that are the accidental byproduct of our evolutionary history? There is no such thing as a biological family, no such thing as human nature or a "normal" human being, and even if there were, why should that be a guide for what is right and just? Huxley is telling us, in effect, that we should continue to feel pain, be depressed or lonely, or suffer from debilitating disease, all because that is what human beings have done for most of their existence as a species. Certainly, no one ever got elected to Congress on such a platform. Instead of taking these characteristics and saying that they are the basis for "human dignity," why don't we simply accept our destiny as creatures who modify themselves?

Huxley suggests that one source for a definition of what it means to be a human being is religion. In *Brave New World*, religion has been abolished and Christianity is a distant memory. The Christian tradition maintains that man is created in God's image, which is the source of human dignity. To use biotechnology to engage in what another Christian writer, C. S. Lewis, called the "abolition of man" is thus a violation of God's will. But I don't think that a careful reading of Huxley or Lewis leads to the conclusion that either writer believed religion to be the *only* grounds on which one could understand the meaning of being human. Both writers suggest that nature itself, and in particular human nature, has a special role in defining for us what is right and wrong, just and unjust, important and unimportant. So our final judgment on "what's wrong" with Huxley's *Brave New World* stands or falls with our view of how important human nature is as a source of values.

The aim of this book is to argue that Huxley was right, that the most significant threat posed by contemporary biotechnology is the possibility that it will alter human nature and thereby move us into a "posthuman" stage of history. This is important, I will argue, because human nature exists, is a meaningful concept, and has provided a stable continuity to our experience as a species. It is, conjointly with religion, what defines our most basic values. Human nature shapes and constrains the possible kinds of political regimes, so a technology powerful enough to reshape what we are will have possibly malign consequences for liberal democracy and the nature of politics itself.

It may be that, as in the case of *1984*, we will eventually find biotechnology's consequences are completely and surprisingly benign, and that we were wrong to lose sleep over it. It may be that the technology will in the end prove much less powerful than it seems today, or that people will be moderate and careful in their application of it. But one of the reasons I am not quite so sanguine is that biotechnology, in contrast to many other scientific advances, mixes obvious benefits with subtle harms in one seamless package.

Nuclear weapons and nuclear energy were perceived as dangerous from the start, and therefore were subject to strict regulation from the moment the Manhattan Project created the first atomic bomb in 1945. Observers like Bill Joy have worried about nanotechnology—that is, molecular-scale self-replicating machines capable of reproducing out of control and destroying their creators. But such threats are actually the easiest to deal with because they are so obvious. If you are likely to be killed by a machine you've created, you take measures to protect yourself. And so far we've had a reasonable record in keeping our machines under control.

There may be products of biotechnology that will be similarly obvious in the dangers they pose to mankind—for example, superbugs, new viruses, or genetically modified foods that produce toxic reactions. Like nuclear weapons or nanotechnology, these are in a way the easiest to deal with because once we have identified them as dangerous, we can treat them as a straightforward threat. The more typical threats raised by biotechnology, on the other hand, are those captured so well by Huxley, and are summed up in the title of an article by novelist Tom Wolfe, "Sorry, but Your Soul Just Died." Medical technology offers us in many cases a devil's bargain: longer life, but with reduced mental capacity; freedom from depression, together with freedom from creativity or spirit; therapies that blur the line between what

we achieve on our own and what we achieve because of the levels of various chemicals in our brains.

Consider the following three scenarios, all of which are distinct possibilities that may unfold over the next generation or two.

The first has to do with new drugs. As a result of advances in neuropharmacology, psychologists discover that human personality is much more plastic than formerly believed. It is already the case that psychotropic drugs such as Prozac and Ritalin can affect traits like self-esteem and the ability to concentrate, but they tend to produce a host of unwanted side effects and hence are shunned except in cases of clear therapeutic need. But in the future, knowledge of genomics permits pharmaceutical companies to tailor drugs very specifically to the genetic profiles of individual patients and greatly minimize unintended side effects. Stolid people can become vivacious; introspective ones extroverted; you can adopt one personality on Wednesday and another for the weekend. There is no longer any excuse for anyone to be depressed or unhappy; even "normally" happy people can make themselves happier without worries of addiction, hangovers, or long-term brain damage.

In the second scenario, advances in stem cell research allow scientists to regenerate virtually any tissue in the body, such that life expectancies are pushed well above 100 years. If you need a new heart or liver, you just grow one inside the chest cavity of a pig or cow; brain damage from Alzheimer's and stroke can be reversed. The only problem is that there are many subtle and some not-so-subtle aspects of human aging that the biotech industry hasn't quite figured out how to fix: people grow mentally rigid and increasingly fixed in their views as they age, and try as they might, they can't make themselves sexually attractive to each other and continue to long for partners of reproductive age. Worst of all, they just refuse to get out of the way, not just of their children, but their grandchildren and great-grandchildren. On the other hand, so few people have children or any connection with traditional reproduction that it scarcely seems to matter.

In a third scenario, the wealthy routinely screen embryos before implantation so as to optimize the kind of children they have. You can increasingly tell the social background of a young person by his or her looks and intelligence; if someone doesn't live up to social expectations, he tends to blame bad genetic choices by his parents rather than himself. Human genes have been transferred to animals and even to plants, for research purposes and to produce new medical products; and animal genes have been added to certain embryos to increase their physical endurance or resistance to disease. Scientists have not dared to produce a full-scale chimera, half human and half ape, though they could; but young people begin to suspect that classmates who do much less well than they do are in fact genetically not fully human. Because, in fact, they aren't.

Sorry, but your soul just died ...

Toward the very end of his life, Thomas Jefferson wrote, "The general spread of the light of science has already laid open to every view the palpable truth, that the mass of mankind has not been born with saddles on their backs, nor a favored few booted and spurred, ready to ride them legitimately, by the grace of God." The political equality enshrined in the Declaration of Independence rests on the empirical fact of natural human equality. We vary greatly as individuals and by culture, but we share a common humanity that allows every human being to potentially communicate with and enter into a moral relationship with every other human being on the planet. The ultimate question raised by biotechnology is, What will happen to political rights once we are able to, in effect, breed some people with saddles on their backs, and others with boots and spurs?

A STRAIGHTFORWARD SOLUTION

What should we do in response to biotechnology that in the future will mix great potential benefits with threats that are either physical and overt or spiritual and subtle? The answer is obvious: *We should use the power of the state to regulate it.* And if this proves to be beyond the power of any

individual nation-state to regulate, it needs to be regulated on an international basis. We need to start thinking concretely now about how to build institutions that can discriminate between good and bad uses of biotechnology, and effectively enforce these rules both nationally and internationally.

This obvious answer is not obvious to many of the participants in the current biotechnology debate. The discussion remains mired at a relatively abstract level about the ethics of procedures like cloning or stem cell research, and divided into one camp that would like to permit everything and another camp that would like to ban wide areas of research and practice. The broader debate is of course an important one, but events are moving so rapidly that we will soon need more practical guidance on how we can direct future developments so that the technology remains man's servant rather than his master. Since it seems very unlikely that we will either permit everything or ban research that is highly promising, we need to find a middle ground.

The creation of new regulatory institutions is not something that should be undertaken lightly, given the inefficiencies that surround all efforts at regulation. For the past three decades, there has been a commendable worldwide movement to deregulate large sectors of every nation's economy, from airlines to telecommunications, and more broadly to reduce the size and scope of government. The global economy that has emerged as a result is a far more efficient generator of wealth and technological innovation. Excessive regulation in the past has led many to become instinctively hostile to state intervention in any form, and it is this knee-jerk aversion to regulation that will be one of the chief obstacles to getting human biotechnology under political control.

But it is important to discriminate: what works for one sector of the economy will not work for another. Information technology, for example, produces many social benefits and relatively few harms and therefore has appropriately gotten by with a fairly minimal degree of government regulation. Nuclear materials and toxic waste, on the other hand, are subject to strict national and international controls because unregulated trade in them would clearly be dangerous.

One of the biggest problems in making the case for regulating human biotechnology is the common view that even if it were desirable to stop technological advance, it is impossible to do so. If the United States or any other single country tries to ban human cloning or germ-line genetic engineering or any other procedure, people who wanted to do these things would simply move to a more favorable jurisdiction where they were permitted. Globalization and international competition in biomedical research ensure that countries that hobble themselves by putting ethical constraints on their scientific communities or biotechnology industries will be punished.

The idea that it is impossible to stop or control the advance of technology is simply wrong.... We in fact control all sorts of technologies and many types of scientific research; people are no more free to experiment in the development of new biological warfare agents than they are to experiment on human subjects without the latter's informed consent. The fact that there are some individuals or organizations that violate these rules, or that there are countries where the rules are either nonexistent or poorly enforced, is no excuse for not making the rules in the first place. People get away with robbery and murder, after all, which is not a reason to legalize theft and homicide.

We need at all costs to avoid a defeatist attitude with regard to technology that says that since we can't do anything to stop or shape developments we don't like, we shouldn't bother trying in the first place. Putting in place a regulatory system that would permit societies to control human biotechnology will not be easy: it will require legislators in countries around the world to step up to the plate and make difficult decisions on complex scientific issues. The shape and form of the institutions designed to implement new rules is a wide-open question; designing them to be minimally obstructive of positive developments while giving them effective enforcement capabilities is a significant challenge. Even more challenging will be the creation of common rules at

an international level, the forging of a consensus among countries with different cultures and views on the underlying ethical questions. But political tasks of comparable complexity have been successfully undertaken in the past.

BIOTECHNOLOGY AND THE RECOMMENCEMENT OF HISTORY

Many of the current debates over biotechnology, on issues like cloning, stem cell research, and germ-line engineering, are polarized between the scientific community and those with religious commitments. I believe that this polarization is unfortunate because it leads many to believe that the *only* reason one might object to certain advances in biotechnology is out of religious belief. Particularly in the United States, biotechnology has been drawn into the debate over abortion; many researchers feel that valuable progress is being checked out of deference to a small number of antiabortion fanatics.

I believe that it is important to be wary of certain innovations in biotechnology for reasons that have nothing to do with religion. The case that I will lay out here might be called Aristotelian, not because I am appealing to Aristotle's authority as a philosopher, but because I take his mode of rational philosophical argument about politics and nature as a model for what I hope to accomplish.

Aristotle argued, in effect, that human notions of right and wrong—what we today call human rights—were ultimately based on human nature. That is, without understanding how natural desires, purposes, traits, and behaviors fit together into a human whole, we cannot understand human ends or make judgments about right and wrong, good and bad, just and unjust. Like many more recent utilitarian philosophers, Aristotle believed that the good was defined by what people desired; but while utilitarians seek to reduce human ends to a simple common denominator like the relief of suffering or the maximization of pleasure, Aristotle retained a complex and nuanced view of the diversity and greatness of natural human ends. The purpose of his philosophy was to try to differentiate the natural from the conventional, and to rationally order human goods.

Aristotle, together with his immediate predecessors Socrates and Plato, initiated a dialog about the nature of human nature that continued in the Western philosophical tradition right up to the early modern period, when liberal democracy was born. While there were significant disputes over what human nature was, no one contested its importance as a basis for rights and justice. Among the believers in natural right were the American Founding Fathers, who based their revolution against the British crown on it. Nonetheless, the concept has been out of favor for the past century or two among academic philosophers and intellectuals.

... I believe this is a mistake, and that any meaningful definition of rights must be based on substantive judgments about human nature. Modern biology is finally giving some meaningful empirical content to the concept of human nature, just as the biotech revolution threatens to take the punch bowl away.

Whatever academic philosophers and social scientists may think of the concept of human nature, the fact that there has been a stable human nature throughout human history has had very great political consequences. As Aristotle and every serious theorist of human nature has understood, human beings are by nature cultural animals, which means that they can learn from experience and pass on that learning to their descendants through nongenetic means. Hence human nature is not narrowly determinative of human behavior but leads to a huge variance in the way people raise children, govern themselves, provide resources, and the like. Mankind's constant efforts at cultural self-modification are what lead to human history and to the progressive growth in the complexity and sophistication of human institutions over time.

The fact of progress and cultural evolution led many modern thinkers to believe that human beings were almost infinitely plastic—that is, that they could be shaped by their social environment to behave in open-ended ways. It is here that

the contemporary prejudice against the concept of human nature starts. Many of those who believed in the social construction of human behavior had strong ulterior motives: they hoped to use social engineering to create societies that were just or fair according to some abstract ideological principle. Beginning with the French Revolution, the world has been convulsed with a series of Utopian political movements that sought to create an earthly heaven by radically rearranging the most basic institutions of society, from the family to private property to the state. These movements crested in the twentieth century, with the socialist revolutions that took place in Russia, China, Cuba, Cambodia, and elsewhere.

By the end of the century, virtually every one of these experiments had failed, and in their place came efforts to create or restore equally modern but less politically radical liberal democracies. One important reason for this worldwide convergence on liberal democracy had to do with the tenacity of human nature. For while human behavior is plastic and variable, it is not infinitely so; at a certain point deeply rooted natural instincts and patterns of behavior reassert themselves to undermine the social engineer's best-laid plans. Many socialist regimes abolished private property, weakened the family, and demanded that people be altruistic to mankind in general rather than to a narrower circle of friends and family. But evolution did not shape human beings in this fashion. Individuals in socialist societies resisted the new institutions at every turn, and when socialism collapsed after the fall of the Berlin Wall in 1989, older, more familiar patterns of behavior reasserted themselves everywhere.

Political institutions cannot abolish either nature or nurture altogether and succeed. The history of the twentieth century was defined by two opposite horrors, the Nazi regime, which said biology was everything, and communism, which maintained that it counted for next to nothing. Liberal democracy has emerged as the only viable and legitimate political system for modern societies because it avoids either extreme, shaping politics according to historically created norms of justice while not interfering excessively

There were many other factors affecting the trajectory of history, which I discussed in my book *The End of History and the Last Man*. One of the basic drivers of the human historical process has been the development of science and technology, which is what determines the horizon of economic production possibilities and therefore a great deal of a society's structural characteristics. The development of technology in the late twentieth century was particularly conducive to liberal democracy. This is not because technology promotes political freedom and equality per se—it does not—but because late-twentieth-century technologies (particularly those related to information) are what political scientist Ithiel de Sola Pool has labeled technologies of freedom.

There is no guarantee, however, that technology will always produce such positive political results. Many technological advances of the past reduced human freedom. The development of agriculture, for example, led to the emergence of large hierarchical societies and made slavery more feasible than it had been in hunter-gatherer times. Closer to our own time, Eli Whitney's invention of the cotton gin made cotton a significant cash crop in the American South at the beginning of the nineteenth century and led to the revitalization of the institution of slavery there.

As the more perceptive critics of the concept of the "end of history" have pointed out, there can be no end of history without an end of modern natural science and technology. Not only are we not at an end of science and technology; we appear to be poised at the cusp of one of the most momentous periods of technological advance in history. Biotechnology and a greater scientific understanding of the human brain promise to have extremely significant political ramifications. Together, they reopen possibilities for social engineering on which societies, with their twentieth-century technologies, had given up.

If we look back at the tools of the past century's social engineers and Utopian planners, they seem unbelievably crude and unscientific. Agitprop, labor camps, reeducation, Freudianism, early childhood conditioning, behavioralism—all

peg of human nature into the round hole of social planning. None of them were based on knowledge of the neurological structure or biochemical basis of the brain; none understood the genetic sources of behavior, or if they did, none could do anything to affect them.

All of this may change in the next generation or two. We do not have to posit a return of state-sponsored eugenics or widespread genetic engineering to see how this could happen. Neuropharmacology has already produced not just Prozac for depression but Ritalin to control the unruly behavior of young children. As we discover not just correlations but actual molecular pathways between genes and traits like intelligence, aggression, sexual identity, criminality, alcoholism, and the like, it will inevitably occur to people that they can make use of this knowledge for particular social ends. This will play itself out as a series of ethical questions facing individual parents, and also as a political issue that may someday come to dominate politics. If wealthy parents suddenly have open to them the opportunity to increase the intelligence of their children as well as that of all their subsequent descendants, then we have the makings not just of a moral dilemma but of a full-scale class war.

THE POLITICAL CONTROL OF BIOTECHNOLOGY

... Holy cruelty.—A man who held a newborn child in his hands approached a holy man. "What shall I do with this child?" he asked; "it is wretched, misshapen, and does not have life enough to die." "Kill it!" shouted the holy man with a terrible voice; "and then hold it in your arms for three days and three nights to create a memory for yourself. Never again will you beget a child this way when it is not time for you to beget."—When the man had heard this, he walked away, disappointed, and many people reproached the holy man because he had counseled cruelty; for he had counseled the man to kill the child. "But is it not crueler to let it live?" asked the holy man.

FRIEDRICH NIETZSCHE,
THE GAY SCIENCE, SECTION 731

Some new technologies are frightening from the start and create an instant consensus on the need to establish political controls over their development and use. When the first atomic bomb was detonated at Alamogordo, New Mexico, in the summer of 1945, not one of the witnesses to the event failed to understand that a terrible new potential for destruction had been created. Nuclear weapons were thus from the start ringed with political controls: individuals could not freely develop nuclear technology on their own or traffic in the parts necessary to create atomic bombs, and, in time, nations that became signatories to the 1968 nonproliferation treaty agreed to control international trade in nuclear technology.

Other new technologies appear to be much more benign, and consequently subject to little or no regulation. Personal computers and the Internet are examples: these new forms of information technology (IT) promised to create wealth, spread access to information and therefore power around more democratically, and foster community among their users. People had to look hard for downsides to the Information Revolution; what they have found to date are issues like the so-called digital divide (that is, inequality of access to IT) and threats to privacy, neither of which qualify as earth-shaking matters of justice or morality. Despite occasional efforts on the part of the world's more statist societies to try to control the use of IT, it has blossomed in recent years with minimal regulatory oversight on either a national or international level.

Biotechnology falls somewhere between these extremes. Transgenic crops and human genetic engineering make people far more uneasy than do personal computers and the Internet. But biotechnology also promises important benefits for human health and well-being. When presented with an advance like the ability to cure a child of cystic fibrosis or diabetes, people find it difficult to articulate reasons why their unease with the technology should stand in the way of progress. It is easiest to object to a new biotechnology if its development leads to a botched clinical trial or to a deadly allergic reaction to a genetically modified food. But the real threat of biotechnology is

far more subtle, and therefore harder to weigh in any utilitarian calculus.

In the face of the challenge from a technology like this, where good and bad are intimately connected, it seems to me that there can be only one possible response: countries must regulate the development and use of technology politically, setting up institutions that will discriminate between those technological advances that promote human flourishing, and those that pose a threat to human dignity and well-being. These regulatory institutions must first be empowered to enforce these discriminations on a national level, and must ultimately extend their reach internationally.

The state of the debate on biotechnology is today polarized between two camps. The first is libertarian, and argues that society should not and cannot put constraints on the development of new technology. This camp includes researchers and scientists who want to push back the frontiers of science, the biotech industry that stands to profit from unfettered technological advance, and, particularly in the United States and Britain, a large group that is ideologically committed to some combination of free markets, deregulation, and minimal government interference in technology.

The other camp is a heterogeneous group with moral concerns about biotechnology, consisting of those who have religious convictions, environmentalists with a belief in the sanctity of nature, opponents of new technology, and people on the Left who are worried about the possible return of eugenics. This group, which ranges from activists like Jeremy Rifkin to the Catholic Church, has proposed banning a wide array of new technologies, from in vitro fertilization and stem cell research to transgenic crops and human cloning.

The debate on biotechnology has to move beyond this polarization. Both approaches—a totally laissez-faire attitude toward biotech development, and the attempt to ban wide swaths of future technology—are misguided and unrealistic. Certain technologies, such as human cloning, deserve to be banned outright, for reasons both intrinsic and tactical. But for most other forms of biotechnology that we see emerging, a more nuanced regulatory approach will be needed. While everyone has been busy staking out ethical positions pro and con various technologies, almost no one has been looking concretely at what kinds of institutions would be needed to allow societies to control the pace and scope of technology development.

It has been a long time since anyone has proposed that what the world needs is more regulation. Regulation—and particularly international regulation—is not something that anyone should call for lightly. Before the Reagan-Thatcher revolutions of the 1980s, many sectors of the economies of North America, Europe, and Japan were vastly over-regulated, and many continue to be so today. Regulation brings with it many inefficiencies and even pathologies that are well understood. Research has shown, for example, how government regulators develop a self-interest in promoting their own power and position, even as they make claims to speak in the public interest. Poorly thought out regulation can drive up the costs of doing business enormously, stifle innovation, and lead to the misallocation of resources as businesses seek to avoid burdensome rules. A great deal of innovative work has been done in the past generation on alternatives to formal state regulation—for example, the self-regulation of businesses, and more flexible models for rule generation and enforcement.

The inefficiency of any scheme of regulation is a fact of life. We can try to minimize it by designing institutions that seek to streamline the regulatory process and make it more responsive to changes in technology and social needs, but in the end there are certain types of social problems that can only be addressed through formal government control. Schemes for self-regulation tend to work best in situations in which an industry doesn't produce a lot of social costs (negative externalities, in economists' terminology), in which the issues tend to be technical and apolitical, and in which industry itself has strong incentives to police itself. This is true in international standards setting, coordination of airline traffic routes and payments, product testing, and

bank settlements and was at one time true for food safety and medical experimentation.

But it is not true of present-day biotechnology, or of the kinds of biomedical technologies that are likely to appear in the future. While the community of research scientists has in the past done an admirable job in policing itself in such areas as human experimentation and the safety of recombinant DNA technology, there are now too many commercial interests chasing too much money for self-regulation to continue to work well into the future. Most biotechnology companies will simply not have the incentives to observe many of the fine ethical distinctions that need to be made, which means that governments necessarily have to step in to draw up and enforce rules for them.

Many people today believe that biotechnology should not and cannot, as a practical matter, be controlled. Both of these conclusions are wrong, as we will see.

WHO GETS TO DECIDE?

Who gets to decide whether we will control a new biotechnology, and with what authority?

During the debate in the U.S. Congress on bills to ban human cloning in 2001, Congressman Ted Strickland of Ohio insisted that we be guided strictly by the best available science, and that "We should not allow theology, philosophy, or politics to interfere with the decision we make on this issue."

There are many who would agree with this statement. Opinion polls in most countries show the public holding scientists in much higher regard than politicians, not to mention theologians or philosophers. Legislators, as we well know, like to posture, exaggerate, argue by anecdote, pound the table, and pander. They often speak and act out of ignorance and are at times heavily influenced by lobbyists and entrenched interests. Why should they, rather than the disinterested community of researchers, have final say on highly complex and technical issues like biotechnology? Efforts by politicians to limit what scientists do in their own domain evoke memories of the medieval Catholic Church branding Galileo a heretic for saying the earth revolves around the sun. Since the time of Francis Bacon, the pursuit of scientific research has been seen to carry its own legitimacy as an activity that automatically serves the broader interests of mankind.

This view is, unfortunately, not correct.

Science by itself cannot establish the ends to which it is put. Science can discover vaccines and cures for diseases, but it can also create infectious agents; it can uncover the physics of semiconductors but also the physics of the hydrogen bomb. Science qua science is indifferent to whether data are gathered under rules that scrupulously protect the interests of human research subjects. Data, after all, are data, and better data can often be obtained … by bending the rules or ignoring them altogether. A number of the Nazi doctors who injected concentration camp victims with infectious agents or tortured prisoners by freezing or burning them to death were in fact legitimate scientists who gathered real data that could potentially be put to good use.

It is only "theology, philosophy, or politics" that can establish the ends of science and the technology that science produces, and pronounce on whether those ends are good or bad. Scientists may help establish moral rules concerning their own conduct, but they do so not as scientists but as scientifically informed members of a broader political community. There are very many brilliant, dedicated, energetic, ethical, and thoughtful people within the community of research scientists and doctors working in the field of biomedicine. But their interests do not necessarily correspond to the public interest. Scientists are strongly driven by ambition, and often have pecuniary interests in a particular technology or medicine as well. Hence the question of what we do with biotechnology is a political issue that cannot be decided technocratically.

The answer to the question of who gets to decide on the legitimate and illegitimate uses of science is actually pretty simple, and has been established by several centuries of political theory and practice: it is the democratically constituted political community, acting chiefly through their elected representatives, that is sovereign in these

matters and has the authority to control the pace and scope of technological development, While there are all sorts of problems with democratic institutions today, from special-interest lobbying to populist posturing, there is also no obviously better alternative set of institutions that can capture the will of the people in a fair and legitimate way. We can surely hope that politicians make decisions that are informed by a sophisticated understanding of science. History is full of cases where laws were made based on bad science, such as the eugenics legislation passed in the United States and Europe in the early twentieth century. But in the end, science itself is just a tool for achieving human ends; what the political community decides are appropriate ends are not ultimately scientific questions.

When we turn to the question of establishing a regulatory regime for human biotechnology, we face a rather different problem. The issue is not whether it should be scientists or politicians who make choices regarding scientific research, but whether what is best in terms of reproductive decisions should be decided by individual parents or the government. James Watson has argued that it should be individual mothers and not a group of male regulators:

> My principle here is pretty simple: just have most of the decisions made by women as opposed to men. They're the ones who bear children, and men, as you know, often sneak away from children that aren't healthy. We're going to have to feel more responsible for the next generation. I think women should be allowed to make the decisions, and as far as I'm concerned, keep these male doctor committees out of action.

Counterpoising the judgment of male bureaucrats against the concerns of loving mothers is a clever rhetorical strategy, but it is beside the point. Male judges, officers, and social workers (as well as a lot of female ones) already interfere in the lives of women all the time, telling them they can't neglect or abuse their children, that they have to send them to school rather than making them go out to earn money for the family, or that they can't give them drugs or arm them with weapons. The fact that most women will use their authority responsibly doesn't eliminate the need for rules, particularly when technology makes possible all sorts of highly unnatural reproductive possibilities (like cloning) whose ultimate consequences for children may not be healthy.

... The automatic community of interest that is assumed to exist between parent and child under natural forms of reproduction may not exist under the new ones. Some have argued that we can presume the consent of a yet-to-be-born child to be free of birth defects or of mental retardation. But can we presume the consent of a child to be a clone, or to be born the biological offspring of two women, or to be born with a nonhuman gene? Cloning in particular raises the prospect that the reproductive decision will suit the interests and convenience of the parent rather than those of the child, and in this case, the state has an obligation to intervene to protect the child.

CAN TECHNOLOGY BE CONTROLLED?

Even if we decide that technology should legitimately be controlled, we face the problem of whether it can be. Indeed, one of the greatest obstacles to thinking about a regulatory scheme for human biotechnology is the widespread belief that technological advance cannot be regulated, and that all such efforts are self-defeating and doomed to failure. This is asserted gleefully by enthusiasts of particular technologies and by those who hope to profit from them, and pessimistically by those who would like to slow down the spread of potentially harmful technologies. In the latter camp, particularly, there is a kind of defeatism as to the ability of politics to shape the future.

This belief has become particularly strong in recent years because of the advent of globalization and our recent experience with information technology. No sovereign nation-state, it is said, can regulate or ban any technological innovation, because the research and development will simply move to another jurisdiction. American efforts to control data encryption, for example, or French

French Web sites, have simply hobbled technological development in these countries, as developers moved their operations to more favorable regulatory climates. The only way to control the spread of technology is to have international agreements on technology-restricting rules, which are extraordinarily difficult to negotiate and even harder to enforce. In the absence of such international agreements, any nation that chooses to regulate itself will simply give other nations a leg up.

This kind of pessimism about the inevitability of technological advance is wrong, and it could become a self-fulfilling prophecy if believed by too many people. For it is simply not the case that the speed and scope of technological development cannot be controlled. There are many dangerous or ethically controversial technologies that have been subject to effective political control, including nuclear weapons and nuclear power, ballistic missiles, biological and chemical warfare agents, replacement human body parts, neuropharmacological drugs, and the like, which cannot be freely developed or traded internationally. The international community has regulated human experimentation effectively for many years. More recently, the proliferation of genetically modified organisms (GMOs) in the food chain has been stopped dead in its tracks in Europe, with American farmers walking away from transgenic crops that they had only recently embraced. One can argue about the rightness of this decision on scientific grounds, but it proves that the march of biotechnology is not an unstoppable juggernaut.

Indeed, the common assumption that it is impossible to control pornography or political discussion on the Internet is wrong. It is not possible for a government to shut down every objectionable Web site around the world, but it is possible to raise the costs of accessing them for ordinary people who live in their jurisdictions. The Chinese authorities, for instance, have used their political power effectively to force Internet companies like Yahoo! and MSN to restrict publication of unsympathetic stories on their Chinese-language Web sites by simply threatening

Skeptics will argue that none of these efforts to control technology has been successful in the end. Despite the huge diplomatic effort that the West, and especially the United States, has put into nonproliferation, for example, India and Pakistan became the sixth and seventh powers to test nuclear devices openly in the 1990s. While nuclear power for energy generation was slowed down after Three Mile Island and Chernobyl, it is now back on the table because of rising fossil fuel costs and concerns over global warming. Ballistic missile proliferation and the development of weapons of mass destruction continue in places like Iraq and North Korea, while there is a large underground market in drugs, spare body parts, plutonium, and virtually any other illicit commodity one cares to name.

All of this is true enough: no regulatory regime is ever fully leak-proof, and if one selects a sufficiently long time frame, most technologies end up being developed eventually. But this misses the point of social regulation: no law is ever fully enforced. Every country makes murder a crime and attaches severe penalties to homicide, and yet murders nonetheless occur. The fact that they do has never been a reason for giving up on the law or on attempts to enforce it.

In the case of nuclear weapons, vigorous nonproliferation efforts on the part of the international community were in fact very successful in slowing down their spread and keeping them out of the hands of countries that might at different points in their histories have been tempted to use them. At the dawn of the nuclear era, in the late 1940s, experts routinely predicted that dozens of countries would possess nuclear weapons in a few years. That only a handful have developed them, and that none of these weapons had been detonated in conflict by the end of the twentieth century, was a remarkable achievement. There are any number of countries that could have developed nuclear weapons but refrained from doing so. Brazil and Argentina, for example, harbored nuclear ambitions when they were both military dictatorships. The nonproliferation regime in which they were enmeshed, however, forced them to keep these programs secret and

to democracy in the 1980s, the programs were shut down.

Nuclear weapons are easier to control than biotechnology for two reasons. First, nuclear weapons development is very expensive and requires large, visible institutions, making their private development very unlikely. Second, the technology is so obviously dangerous that there was a rapid worldwide consensus on the need to control it. Biotechnology, by contrast, can be carried out in smaller, less lavishly funded labs, and there is no similar consensus on its downside risks.

On the other hand, biotechnology does not pose high enforcement hurdles the way nuclear weapons do. A single bomb in the hands of a terrorist group or rogue state like Iraq will pose significant dangers to the world's security. By contrast, an Iraq that can clone Saddam Hussein does not pose much of a threat, unappetizing as that prospect may be. The purpose of a law banning human cloning in the United States would not be undermined if some other countries in the world permitted it, or if Americans could travel abroad to have themselves cloned in such jurisdictions.

The argument that regulation cannot work in a globalized world unless it is international in scope is true enough, but to use this fact to build a case against national-level regulation is to put the cart before the horse. Regulation seldom starts at an international level: nation-states have to develop rules for their own societies before they can even begin to think about creating an international regulatory system.* This is particularly true in the case of a politically, economically, and culturally dominant country like the United States: other countries around the world will pay a great deal of attention to what the United States does in its domestic law. If an international consensus on the regulation of certain biotechnologies is ever to take shape, it is very difficult to see it coming about in the absence of American action at a national level.

In pointing to other cases where technology has been regulated with some success, I do not mean to underestimate the difficulty of creating a similar system for human biotechnology. The international biotech industry is highly competitive, and companies are constantly searching for the most favorable regulatory climate in which to do their work. Because Germany, with its traumatic history of eugenics, has been more restrictive of genetic research than many developed countries, most German pharmaceutical and biotech companies have moved their labs to Britain, the United States, and other less restrictive countries. In 2000, Britain legalized therapeutic or research cloning and will become a haven for this type of research should the United States join Germany, France, and a number of other countries that do not permit it. Singapore, Israel, and other countries have indicated an interest in pursuing research in stem cells and other niches if the United States continues to restrict its own efforts out of ethical concerns.

The realities of international competition do not mean, however, that the United States or any other country has to fatalistically jump into a technological arms race. We do not know at this point whether an international consensus on the banning or strict regulation of certain technologies, such as cloning and germ-line modification, will emerge, but there is absolutely no reason to rule out the possibility at this early stage in the game.

Take the issue of reproductive cloning—that is, the cloning of a human child. As of this writing (November 2001), twenty-four countries had banned reproductive cloning, including Germany, France, India, Japan, Argentina, Brazil, South Africa, and the United Kingdom. In 1998 the Council of Europe approved an Additional Protocol to its Convention on Human Rights and Dignity with Regard to Bio-medicine, banning human reproductive cloning; the document was approved by twenty-four of the council's forty-three member states. The U.S. Congress

* There are some exceptions to this general rule, such as the case of new or transitional democracies that appeal to international rules on human rights to promote the observance of these rules in their own societies. This analogy is not appropriate in the case of rules for biotechnology, however. International conventions on human rights were established at the instigation of countries that observed these rights and had them codified in their legal systems already.

was just one of a number of other legislatures deliberating on similar measures. The French and German governments have proposed that the United Nations enact a global reproductive cloning ban. Given that Dolly the sheep had been created only four years earlier, it is not surprising that it has taken time for politicians and the law to catch up with technology. But at the moment it appears that much of the world is heading toward a consensus on the illegitimacy of human reproductive cloning. It may be that in a few years, if some crackpot cult like the Raelians wants to clone a child, they will have to travel to North Korea or Iraq to do so.

What are the prospects for the emergence of an international consensus on biotech regulation? It is hard to say at this early point, but it is possible to make some observations about culture and politics with regard to this issue.

There is a continuum of views in the world today concerning the ethicality of certain types of biotechnology and particularly genetic manipulation. At the most restrictive end of this continuum are Germany and other countries in continental Europe that, for historical reasons already mentioned, have been very reluctant to move too far down this road. Continental Europe has also been home to the world's strongest environmental movements, which as a whole have been quite hostile to biotechnology in its various forms.

At the other end of the spectrum are a number of countries in Asia, which for historical and cultural reasons have not been nearly as concerned with the ethical dimension of biotechnology. Much of Asia, for example, lacks religion per se as it is understood in the West—that is, as a system of revealed belief that originates from a transcendental deity. The dominant ethical system in China, Confucianism, lacks any concept of God; folk religions like Taoism and Shinto are animistic and invest both animals and inanimate objects with spiritual qualities; and Buddhism conflates human and natural creation into a single seamless cosmos. Asian traditions such as Buddhism, Taoism, and Shinto tend not to make as sharp an ethical distinction between mankind Christianity. That these traditions perceive a continuity between human and nonhuman nature has allowed them to be, as Frans de Waal points out, more sympathetic to nonhuman animals. But it also implies a somewhat lower degree of regard for the sanctity of human life. Consequently, practices such as abortion and infanticide (particularly female infanticide) have been widespread in many parts of Asia. The Chinese government has permitted practices abhorrent in the West, such as the harvesting of organs from executed prisoners, and passed a eugenics law as recently as 1995.

Between continental Europe and Asia on the continuum lie the English-speaking countries, Latin America, and other parts of the world. America and Britain never developed the phobia for genetic research that Germany and France did, and are by virtue of their liberal traditions more skeptical of state regulation. The United States in particular has always been addicted to technological innovation and, for a host of institutional and cultural reasons, is very good at producing it. The American fondness for technology has been strongly reinforced by the information technology revolution of the last two decades, which has convinced many Americans that technology inevitably promises to be individually liberating and personally enriching. Balanced against this is the strength of conservative religious groups in the United States—Protestant, Catholic, and, increasingly, Muslim—that have up to now acted as brakes on uncontrolled technological advance.

Britain has always been closer to America, with its liberal traditions, than to Germany, but it has paradoxically been home to one of the strongest environmental protest movements opposed to GMOs and agricultural biotechnology. There are probably no deep cultural reasons for this; British skepticism about GMOs is more likely traced to the massive regulatory failure represented by mad cow disease, a failure that has left Britain with the largest population of victims to date of the human form of bovine spongiform encephalopathy (BSE), Creutzfeldt-Jacob disease. BSE has nothing to do with biotechnology, of course, but it did reasonably

credibility of governments that pronounce on the safety of food products. A generation ago, Americans were much more concerned with threats to the environment and eager to regulate them, based on their recent experiences with Love Canal and other environmental disasters.

If there is any region of the world that is likely to opt out of an emerging consensus on the regulation of biotechnology, it is Asia. A number of Asian countries either are not democracies or lack strong domestic constituencies opposed to certain types of biotechnology on moral grounds. Asian countries like Singapore and South Korea have the research infrastructure to compete in biomedicine, and strong economic incentives to gain market share in biotechnology at the expense of Europe and North America. In the future, biotechnology may become an important fracture line in world politics.

An international consensus on the control of new biomedical technologies will not simply spring into being without a great deal of work on the part of the international community and the leading countries within it. There is no magic bullet for creating such a consensus. It will require the traditional tools of diplomacy: rhetoric, persuasion, negotiation, economic and political leverage. But in this respect the problem is not different from the creation of any other international regime, whether in air traffic, telecommunications, nuclear or ballistic missile proliferation, and the like.

The international governance of human biotechnology does not inevitably mean the creation of a new international organization, expanding the United Nations, or setting up an unaccountable bureaucracy. At the simplest level it can come about through the effort of nation-states to harmonize their regulatory policies. For members of the European Union (EU), this harmonization will presumably already have occurred at a European level.

Take, for example, the international regime governing pharmaceuticals. Every industrialized country has a science-based regulatory agency comparable to the American Food and Drug Administration to oversee the safety and effectiveness of drugs. In Britain it is the Medicines Control Agency, in Japan the Pharmaceutical Affairs Council, in Germany the Bundesinstitutfür Arzneimittel und Medizinpro-dukte, and in France the Agence Française du Médicament. The European Community has sought to harmonize the drug-approval process of its member states since 1965 to avoid the duplication and waste involved in filing multiple applications in different national jurisdictions. This led to the establishment in London in 1995 of the European Medicines Evaluation Agency, which was supposed to provide one-stop drug approval shopping at a European level. At the same time, the European Commission convened a multilateral meeting to broaden harmonization beyond Europe (called the International Conference on Harmonization). Although some Americans have criticized this as an effort by Eurocrats to extend their reach to the United States, it remains a voluntary regime that has received strong support from the pharmaceutical industry because it could lead to substantial increases in efficiency.

Before we can discuss how human biotechnology needs to be regulated in the future, however, we need to understand how it is regulated today, and how the current system came into existence. The picture is extraordinarily complex, particularly when seen on an international level, and is one in which the history of agricultural and human biotechnology have been closely intertwined.

HOW BIOTECHNOLOGY IS REGULATED TODAY

There are many different approaches to regulation, ranging from self-regulation by industry or the scientific community with minimal government oversight, to formal regulation by a statutory agency. Formal regulation can, moreover, be more or less intrusive: at one extreme, there can be a close relationship between regulator and regulatee, which often invites charges of industry "capture," but there can also be highly adversarial relationships, in which the regulating agency imposes detailed (and unwanted) rules on the target industry and is subject to constant

litigation. Many of these variants have already been applied to biotechnology.

Take genetic engineering. The development of the underlying technology of recombinant DNA (rDNA), in which genes are spliced (often from one species to another), gave rise to an early and exam Review Boards, now required for federally funded research. Even now, the adequacy of these protections has been criticized; the National Bioethics Advisory Commission issued a report in 2001 urging new federal legislation creating a single, strengthened National Office for Human Research Oversight.

Then as now, scientists pursuing ethically questionable research defended their actions on the grounds that the medical benefits that could be derived from their work outweighed possible harms to the research subjects. They also argued that the scientific community alone was best able to judge the risks and benefits of biomedical research, and resisted the intrusion of federal law into their domain.

Rules regarding human experimentation exist on an international level as well. The basic law is the Nuremberg Code, which established the principle that medical experimentation could be performed on a human subject only with the latter's consent. The code grew out of the revelations of the horrifying experiments performed by Nazi doctors on concentration camp inmates during World War II. It had little effect on actual practice in other countries, however, as a recital of later abuses in the United States indicates, and was resisted by many doctors as being too restrictive of valid research.

The Nuremberg Code was largely superseded by the Helsinki Declaration, adopted by the World Medical Association (the global organization representing national medical associations) in 1964. The Helsinki Declaration establishes a number of principles governing experimentation on human subjects, including informed consent, and was better liked by the international medical profession because it was a matter of self-regulation rather than formal international law. Actual practice among developed nations varies a great deal despite these international rules; Japan, for example, has seen a number of cases in the 1990s in which patients were not informed of their conditions or possible treatments by doctors.

Despite variations in practice and occasional lapses, the case of human experimentation shows that the international community is in fact able to place effective limits on the way in which scientific research is conducted, in ways that balance the need for research against respect for the dignity of research subjects. This is an issue that will need to be revisited frequently in the future.

POLICIES FOR THE FUTURE

Advances in biotechnology have created gaping holes in the existing regime for the regulation of human biomedicine, which legislatures and administrative agencies around the world have been racing to fill. It is not clear, for example, that the rules for human experimentation ... apply to embryos outside the womb. The nature of the players and the flow of money within the biomedical and pharmaceutical communities have also changed, with important implications for any future regulatory system.

One thing is reasonably clear: the time when governments could deal with biotech questions by appointing national commissions that brought scientists together with learned theologians, historians, and bioethicists—groups like the National Bioethics Advisory Commission in the United States and the European Group on Ethics in Science and New Technologies—is rapidly drawing to a close. These commissions played a very useful role in doing the preliminary intellectual spadework of thinking through moral and social implications of biomedical research. But it is time to move from thinking to acting, from recommending to legislating. We need institutions with real enforcement powers.

The community of bioethicists that has grown up in tandem with the biotech industry is in many respects a double-edged sword. On the one hand, it has played an extremely useful function by raising doubts and questions about the wisdom and morality of certain technological innovations. On the other hand, many bioethicists have become nothing more than sophisticated (and sophistic) justifiers of whatever it is the scientific community

wants to do, having enough knowledge of Catholic theology or Kantian metaphysics to beat back criticisms by anyone coming out of these traditions who might object more strenuously. The Human Genome Project from the beginning devoted 3 percent of its budget to studying the Ethical, Social, and Legal Implications of genetic research. This can be regarded as commendable concern for the ethical dimensions of scientific research, or else as a kind of protection money the scientists have to pay to keep the true ethicists off their backs. In any discussion of cloning, stem cell research, germ-line engineering, and the like, it is usually the professional bioethicist who can be relied on to take the most permissive position of anyone in the room.* But if the ethicist isn't going to tell you that you can't do something, who will?

A number of countries have in fact moved beyond the stage of national commissions and study groups to actual legislation. One of the first and most contentious policy issues legislators have tried to grapple with concerns the uses that may be made of human embryos. This issue touches on a whole host of medical practices and procedures, both existing today and yet to be developed. These include abortion, in vitro fertilization, preimplantation diagnosis and screening, sex selection, stem cell research, cloning for reproductive and research purposes, and germ-line engineering. There are a huge number of permutations and combinations of possible rules that societies can establish regarding embryos. For example, one can imagine permitting them to be aborted or discarded by in vitro fertilization clinics, yet not created deliberately for research purposes nor selected for sex or other characteristics. Formulation and enforcement of these rules will constitute the substance of any future regulatory system for human biotechnology. There are at present a wide variety of national-level rules regarding human embryos. To date (November 2001), sixteen countries have passed laws regulating human embryo research, including France, Germany, Austria, Switzerland, Norway, Ireland, Poland, Brazil, and Peru (despite the fact that in France abortion is legal). In addition, Hungary, Costa Rica, and Ecuador implicitly restrict research by conferring on embryos a right to life. Finland, Sweden, and Spain permit embryo research, but only on extra embryos left over from in vitro fertilization clinics. Germany's laws are among the most restrictive; since passage of the 1990 Act for the Protection of Embryos (GesetzzumSchutz von Embryonen), a number of areas have been regulated, including abuse of human embryos, sex selection, artificial modification of human germ-line cells, cloning, and the creation of chimeras and hybrids.

Britain in 1990 passed the Fertilisation and Embryology Act, which established one of the most clear-cut legal frameworks in the world for the regulation of embryo research and cloning. This act was thought to ban reproductive cloning while permitting research cloning, though in 2001 a British court ruled reproductive cloning would actually be permitted under a loophole that the government moved quickly to try to close. Given the lack of consensus across the continent on this issue, there has been no action on the European level to regulate embryo research apart from the creation of the European Group on Ethics in Science and New Technologies.

Embryo research is only the beginning of a series of new developments created by technology for which societies have to decide on rules and regulatory institutions. Others that will come up sooner or later include:

Preimplantation diagnosis and screening. This group of technologies, in which multiple embryos are screened genetically for birth defects and other characteristics, is the beginning point of

* This phenomenon is a common one and is known as regulatory "capture," whereby the group that is supposed to be overseeing the activities of an industry becomes an agent for the industry. This happens for many reasons, including the dependence of the regulators on the regulatees for money and information. In addition, there are the career incentives that most professional bioethicists face. Scientists do not usually have to worry about winning the respect of ethicists, particularly if they are Nobel Prize winners in molecular biology or physiology. On the other hand, ethicists face an uphill struggle winning the respect of the scientists they must deal with, and are hardly likely to do so if they tell them they are morally wrong or if they depart significantly from the materialist worldview that the scientists hold dear.

"designer babies" and will arrive much sooner than human germ-line engineering. Indeed, such screening has already been performed for children of parents susceptible to certain genetic diseases. In the future, do we want to permit parents to screen and selectively implant embryos on the basis of sex; intelligence; looks; hair, eye, or skin color; sexual orientation; and other characteristics once they can be identified genetically?

Germ-line engineering. If and when human germ-line engineering arrives, it will raise the same issues as preimplantation diagnosis and screening, but in a more extreme form. Preimplantation diagnosis and screening is limited by the fact that there will always be a limited number of embryos from which to choose, based on the genes of the two parents. Germ-line engineering will expand the possibilities to include virtually any other genetically governed trait, provided it can be identified successfully, including traits that come from other species.

The creation of chimeras using human genes. Geoffrey Bourne, former director of the Emory University primate center, once stated that "it would be very important scientifically to try to produce an ape-human cross," Other researchers have suggested using women as "hosts" for the embryos of chimpanzees or gorillas. One biotech company, Advanced Cell Technology, reported that it had successfully transferred human DNA into a cow's egg and gotten it to grow into a blastocyst before it was destroyed. Scientists have been deterred from doing research in this area for fear of bad publicity, but in the United States such work is not illegal. Will we permit the creation of hybrid creatures using human genes?

New psychotropic drugs. In the United States, the Food and Drug Administration (FDA) regulates therapeutic drugs, while the Drug Enforcement Administration (DEA) and the states regulate illegal narcotics such as heroin, cocaine, and marijuana. Societies will have to make decisions on the legality and extent of permissible use of future generations of neuropharmacological agents. In the case of prospective drugs that improve memory or other cognitive skills, they will have to decide on the desirability of enhancement uses and how they are to be regulated.

WHERE DO WE DRAW RED LINES?

Regulation is essentially the act of drawing a series of red lines that separate legal from proscribed activities, based on a statute that defines the area in which regulators can exercise some degree of judgment. With the exception of some die-hard libertarians, most people reading the above list of innovations that may be made possible by biotechnology will probably want to see some red lines drawn.

There are certain things that should be banned outright. One of them is reproductive cloning—that is, cloning with the intent of producing a child. The reasons for this are both moral and practical, and go way beyond the National Bioethics Advisory Commission's concerns that human cloning cannot now be done safely.

The moral reasons have to do with the fact that cloning is a highly unnatural form of reproduction that will establish equally unnatural relationships between parents and children. A cloned child will have a very asymmetrical relationship with his or her parents. He or she will be both child and twin of the parent from whom his or her genes come, but will not be related to the other parent in any way. The unrelated parent will be expected to nurture a younger version of his or her spouse. How will that parent look upon the clone when he or she reaches sexual maturity? Nature, for all of the reasons explicated in earlier chapters of this book, is a valid point of reference for our values and should not be discarded as a standard for parent-child relationships lightly. While it is possible to come up with some sympathetic scenarios in which cloning might be justified (for example, a Holocaust survivor with no other way of continuing the family line), they do not constitute a sufficiently strong societal interest to justify a practice that on the whole would be harmful.

Beyond these considerations inherent to cloning itself, there are a number of practical

concerns. Cloning is the opening wedge for a series of new technologies that will ultimately lead to designer babies and one that is likely to become feasible much sooner than genetic engineering. If we get used to cloning in the near term, it will be much harder to oppose germ-line engineering for enhancement purposes in the future. It is important to lay down a political marker at an early point to demonstrate that the development of these technologies is not inevitable, and that societies can take some measure of control over the pace and scope of technological advance. There is no strong constituency in favor of cloning in any country. It is also an area where considerable international consensus exists in opposition to the procedure. Cloning therefore represents an important strategic opportunity to establish the possibility of political control over biotechnology.

But while a broad-brush ban is appropriate in this case, it will not be a good model for the control of future technologies. Preimplantation diagnosis and screening have begun to be used today to ensure the birth of children free of genetic diseases. The same technology can be used for less laudable purposes, such as sex selection. What we need to do in this case is not ban the procedure but regulate it, drawing red lines not around the procedure itself but within its range of possible uses to distinguish between what is legitimate and what is illegitimate.

One obvious way to draw red lines is to distinguish between therapy and enhancement, directing research toward the former while putting restrictions on the latter. The original purpose of medicine is, after all, to heal the sick, not to turn healthy people into gods. We don't want star athletes to be hobbled by bad knees or torn ligaments, but we also don't want them to compete on the basis of who has taken the most steroids. This general principle would allow us to use biotechnologies to, for example, cure genetic diseases like Huntington's chorea or cystic fibrosis, but not to make our children more intelligent or taller.

The distinction between therapy and enhancement has been attacked on the grounds that there is no way to distinguish between the two in theory, and therefore no way of discriminating in practice. There is a long tradition, argued most forcefully in recent years by the French postmodernist thinker Michel Foucault, which maintains that what society considers to be pathology or disease is actually a socially constructed phenomenon in which deviation from some presumed norm is stigmatized. Homosexuality, to take one example, was long considered unnatural and was classified as a psychiatric disorder until the latter part of the twentieth century, when it was depathologized as part of the growing acceptance of gayness in developed societies. Something similar can be said of dwarfism: human heights are distributed normally, and it is not clear at what point in the distribution one becomes a dwarf. If it is legitimate to give growth hormone to a child who is in the bottom 0.5 percentile for height, who's to say that you can't also prescribe it for someone who is in the fifth percentile, or for that matter in the fiftieth? Geneticist Lee Silver makes a similar argument about future genetic engineering, saying that it is impossible to draw a line between therapy and enhancement in an objective manner: "in every case, genetic engineering will be used to add something to a child's genome that didn't exist in the genomes of either of its parents."

While it is the case that certain conditions do not lend themselves to neat distinctions between pathological and normal, it is also true that there is such a thing as health. As Leon Kass has argued, there is a natural functioning to the whole organism that has been determined by the requirements of the species' evolutionary history, one that is not simply an arbitrary social construction. It has often seemed to me that the only people who can argue that there is no difference in principle between disease and health are those who have never been sick: if you have a virus or fracture your leg, you know perfectly well that something is wrong.

And even in the cases where the borderline between sickness and health, therapy and enhancement, is murkier, regulatory agencies are routinely able to make these distinctions in practice. Take the case of Ritalin....The underlying "disease" that Ritalin is supposed to treat,

attention deficit–hyperactivity disorder (ADHD), is most likely not a disease at all but simply the label that we put on people who are in the tail of a normal distribution of behavior related to focus and attention. This is in fact a classic case of the social construction of pathology: ADHD was not even in the medical lexicon a couple of generations ago. There is, correspondingly, no neat line between what one might label the therapeutic and enhancement uses of Ritalin. At one end of the distribution, there are children almost anyone would say are so hyperactive that normal functioning is impossible for them, and it is hard to object to treating them with Ritalin. At the other end of the distribution are children who have no trouble whatsoever concentrating or interacting, for whom taking Ritalin might be an enjoyable experience that would give them a high just like any other amphetamine. But they would be taking the drug for enhancement rather than for therapeutic reasons, and thus most people would want to prevent them from doing so. What makes Ritalin controversial is all the children in the middle, who meet some but not all of the diagnostic criteria specified in the *Diagnostic and Statistical Manual of Mental Disorders* for the disease and who nonetheless are prescribed the drug by their family physician.

If there was ever a case, in other words, where the distinction between pathology and health in diagnosis, and therapy and enhancement in treatment, is ambiguous, it is ADHD and Ritalin. And yet, regulatory agencies *make and enforce this distinction all the time*. The DEA classifies Ritalin as a Schedule II pharmaceutical that can only be taken for therapeutic purposes with a doctor's prescription; it clamps down on Ritalin's recreational (that is to say, enhancement) use as an amphetamine. That the boundary between therapy and enhancement is unclear does not make the distinction meaningless. My own strong feeling is that the drug is overprescribed in the United States and used in situations in which parents and teachers ought to employ more traditional means of engaging children and shaping their characters. But the current regulatory system, for all its faults, is better than a situation in which Ritalin is either banned altogether or else sold over the counter like cough medicine.

Regulators are called on all the time to make complex judgments that cannot be held up to precise theoretical scrutiny. What constitutes a "safe" level of heavy metals in the soil, or sulfur dioxide in the atmosphere? How does a regulator justify pushing down the level of a particular toxin in drinking water from fifty to five parts per million, when he or she is trading off health consequences against compliance costs? These decisions are always controversial, but in a sense they are easier to make in practice than in theory. For in practice, a properly functioning democratic political system allows people with a stake in the regulator's decision to push and shove against one another until a compromise is reached.

Once we agree in principle that we will need a capability to draw red lines, it will not be a fruitful exercise to spend a lot of time arguing precisely where they should be placed. As in other areas of regulation, many of these decisions will have to be made on a trial-and-error basis by administrative agencies, based on knowledge and experience not available to us at present. What is more important is to think about the design of institutions that can make and enforce regulations on, for example, the use of preimplantation diagnosis and screening for therapeutic rather than enhancement purposes, and how those institutions can be extended internationally.

As noted at the beginning of this chapter, action has to begin with legislatures stepping up to the plate and establishing rules and institutions. This is easier said than done: biotechnology is a technically complex and demanding subject, one moreover that is changing every day, with a wide variety of interest groups pulling in different directions. The politics of biotechnology does not fall into familiar political categories; if one is a conservative Republican or a left-wing Social Democrat, it is not immediately obvious how one should vote on a bill to permit so-called therapeutic cloning or stem cell research. For these reasons many legislators would rather duck the issue, hoping it will get resolved in some other way.

But to not act under conditions of rapid technological change is in effect to make a decision

legitimizing that change. If legislators in democratic societies do not face up to their responsibilities, other institutions and actors will make the decisions for them.

This is particularly true given the peculiarities of the American political system. In the past, it has been the case that the courts have stepped into controversial areas of social policy when the legislature failed to act to negotiate acceptable political rules. In the absence of congressional action on an issue like cloning, it is conceivable that the courts at some later point may be tempted or compelled to step into the breach and discover, for example, that human cloning or research on cloning is a constitutionally protected right. This was a very poor approach to the formulation of law and public policy in the past, one that tainted policies, such as the legalization of abortion, that more properly should have been enacted legislatively. On the other hand, if the American people clearly express their will on human cloning through their democratically elected representatives, the courts will be reluctant to thwart their will through discovery of a new right.

If the legislature does act to put further regulatory controls on human biotechnology, it will face large questions concerning the design of the requisite institutions to implement them. The same issue came up for the United States and the European Community in the 1980s when agricultural biotechnology appeared on the scene: Do we use existing regulatory bodies to do the job, or are the new technologies sufficiently different so that entirely new agencies are required? In the American case, the Reagan administration decided that agricultural biotech did not represent a sufficiently radical break with the past to merit regulation on the basis of process rather than product. It therefore decided to leave regulatory authority with existing agencies like the FDA and the Environmental Protection Agency (EPA), on the basis of their statutory authority. The Europeans, by contrast, decided to regulate on the basis of process and therefore had to create new regulatory procedures for handling biotech products.

All countries face similar decisions today concerning human biotechnology. In the United States, it would be possible to leave regulatory authority with existing institutions like the FDA, the NIH, or consultative groups like the Recombinant DNA Advisory Committee (RAC). It is prudent to be conservative in the creation of new regulatory institutions and additional layers of bureaucracy. On the other hand, there are a number of reasons for thinking that we need to establish new institutions to deal with the challenges of the coming biotech revolution. Not to do so would be like trying to use the Interstate Commerce Commission, which was responsible for regulating trucks, to oversee civil aviation when that industry came into being, rather than creating a separate Federal Aviation Administration.

Let us consider first the case of the United States. An initial reason that existing U.S. institutions are probably not sufficient to regulate future human biotechnology is the question of their narrow mandate. Human biotechnology differs substantially from agricultural biotechnology insofar as it raises a host of ethical questions related to human dignity and human rights that are not an issue for GMOs. While people object to genetically engineered crops on ethical grounds, the most vociferous complaints have had to do with their possible negative consequences for human health and their environmental impact. This is precisely what existing regulatory institutions like the FDA, the EPA, and the U.S. Department of Agriculture have been set up to do. They can be criticized for having the wrong standards or for not being sufficiently cautious, but they are not operating outside their regulatory mandate when they take on genetically modified foods.

Let us suppose that Congress legislatively distinguishes between therapeutic and enhancement uses of preimplantation diagnosis and screening. The FDA is not set up to make politically sensitive decisions concerning the point at which selection for characteristics like intelligence and height ceases to be therapeutic and becomes enhancing, or whether these characteristics can be considered therapeutic at all. The FDA can disapprove a procedure only on the grounds of effectiveness and safety, but there will be many safe and

effective procedures that will nonetheless require regulatory scrutiny. The limits of the FDA's mandate are already evident: it has asserted aright to regulate human cloning on the legally questionable grounds that a cloned child constitutes a medical "product" over which it has authority.

One can always try to amend and expand the FDA's charter, but past experience shows that it is very difficult to change the organizational culture of agencies with a long history. Not only will the agency resist taking on new duties, but a shifting mandate will likely mean it will do its old job less well. This implies the need to create a new agency to oversee the approval of new medicines, procedures, and technologies for human health. In addition to having a broader mandate, this new authority would have to have different staffing. It would have to include not just the doctors and scientists who staff the FDA and oversee clinical trials for new drugs, but other societal voices that are prepared to make judgments about the technology's social and ethical implications.

A second reason that existing institutions are probably not sufficient to regulate biotechnology in the future has to do with changes that have taken place in the research community and the biotech/pharmaceutical industry as a whole over the past generation. There was a period up through the early 1990s when virtually all biomedical research in the United States was funded by the NIH or another federal government agency. This meant that the NIH could regulate that research through its own internal rule-making authority, as in the case of rules concerning human experimentation. Government regulators could work hand in glove with committees of scientific insiders, like the RAC, and be reasonably sure that no one in the United States was doing dangerous or ethically questionable research.

None of this holds true any longer. While the federal government remains the largest source of research funding, there is a huge amount of private investment money available to sponsor work in new biotechnologies. The U.S. biotech industry by itself spent nearly $11 billion on research in 2000, employs over 150,000 people, and has doubled in size since 1993. Indeed, the massive government-funded Human Genome Project was upstaged by Craig Venter's privately held Celera Genomics in the race to map the human genome. The first embryonic stem cell lines were cultivated by James Thompson at the University of Wisconsin, using nongovernment funding in order to comply with the ban on federally funded research that would harm embryos. Many of the participants at a workshop held on the twenty-fifth anniversary of the Asilomar Conference on rDNA concluded that while the RAC had served an important function in its day, it could no longer monitor or police the present-day biotech industry. It has no formal enforcement powers and can bring to bear only the weight of opinion within the elite scientific community. The nature of that community has changed over time as well: there are today many fewer "pure" researchers, with no ties to the biotech industry or commercial interests in certain technologies.

This means that any new regulatory agency not only would have to have a mandate to regulate biotechnology on grounds broader than efficacy and safety but also would have to have statutory authority over all research and development, and not just research that is federally funded. Such an agency, the Human Fertilisation and Embryology Authority, has already been created in Britain for this purpose. Unification of regulatory powers into a single new agency will end the practice of complying with federal funding restrictions by finding private sponsors and, it is hoped, will shed a more uniform light on the whole biotech sector.

What are the prospects for the United States and other countries putting into place a regulatory system of the kind just outlined? There will be formidable political obstacles to creating new institutions. The biotech industry is strongly opposed to regulation (if anything, it would like to see FDA rules loosened), as is, by and large, the community of research scientists. Most would prefer regulation to take place within their own communities, outside the scope of formal law. They are joined in this by advocacy groups representing patients, the elderly, and others

with an interest in promoting cures for various diseases, and together these groups form a very powerful political coalition.

There are reasons the biotech industry should consider actively promoting the right kind of formal regulation of human biotechnology, however, out of simple long-term self-interest. For that, it needs to look no further than what happened to agricultural biotechnology, which is a good object lesson in the political pitfalls of advancing a new technology too quickly.

At the beginning of the 1990s, Monsanto, a leading innovator in agricultural biotechnology, considered asking the first Bush administration for stronger formal regulatory rules for its genetically engineered products, including labeling requirements. A change of leadership at the top scuttled this initiative, however, on the grounds that there was no scientific evidence of health risks, and the firm introduced a series of new GMOs that were quickly adopted by American farmers. What the company failed to anticipate was the political backlash that arose in Europe against GMOs, and the strict labeling requirements that the European Union imposed in 1997 for genetically modified foods imported into Europe.

Monsanto and other American firms railed at the Europeans for being unscientific and protectionist, but Europe had sufficient market power to enforce its rules on American exporters. American farmers, without a means of separating genetically modified from nongenetically modified foods, found themselves closed out of important export markets. They responded by planting fewer genetically modified crops after 1997 and charging that they had been misled by the biotech industry. In retrospect, Monsanto executives realized that they had made a big mistake by not working earlier to establish an acceptable regulatory environment that would assure consumers of the safety of their products, even if this did not appear to be scientifically necessary.

The history of pharmaceutical regulation was driven by horror stories like the sulfanilamide elixir and thalidomide. It may be the case that regulations concerning human cloning will have to await the birth of a horribly deformed child who is the product of an unsuccessful cloning attempt. The biotech industry needs to consider whether it is better to anticipate such problems now and work toward formulating a system that serves its interests by assuring people of the safety and ethical nature of its products, or wait until there is a huge public outcry following an outrageous accident or horrifying experiment.

THE BEGINNING OF POSTHUMAN HISTORY?

The American regime was built, beginning in 1776, on a foundation of natural right. Constitutional government and a rule of law, by limiting the arbitrary authority of tyrants, would protect the kind of freedom that human beings by nature enjoyed. As Abraham Lincoln pointed out eighty-seven years later, it was also a regime dedicated to the proposition that all men are created equal. There would be an equality of freedom only because there was a natural equality of man; or, to put it more positively, the fact of natural equality demanded an equality of political rights.

Critics have pointed out that the United States has never lived up to this ideal of an equality of freedom and has, in its history, excluded entire groups from it. Defenders of the American regime have, more correctly in my view, pointed out that the principle of equal rights has driven a steady expansion in the circle of those entitled to rights. Once it was established that all human beings have natural rights, the big arguments in American political history have been over who falls within that charmed circles of "men" who were said by the Declaration to be created equal. The circle did not initially encompass women, or blacks, or white men without property; however, it was slowly but surely expanded to encompass them in time.

Whether the participants in these arguments recognized it or not, they all had at least an implicit idea of what the "essence" of a human being was and therefore a ground for judging whether one or another individual

speak, and act very differently from one another, so much of this argument revolved around the question of whether those apparent differences were ones of convention only, or whether they were rooted in nature. Modern natural science has cooperated to some extent in expanding our view of who qualifies as a human being because it has tended to show that most of the apparent differences between human beings are conventional rather than natural. Where there *are* natural differences, as between men and women, they have been shown to affect nonessential qualities that do not have a bearing on political rights.

So, despite the poor repute in which concepts like natural rights are held by academic philosophers, much of our political world rests on the existence of a stable human "essence" with which we are endowed by nature, or rather, on the fact that we believe such an essence exists.

We may be about to enter into a posthuman future, in which technology will give us the capacity gradually to alter that essence over time. Many embrace this power, under the banner of human freedom. They want to maximize the freedom of parents to choose the kind of children they have, the freedom of scientists to pursue research, and the freedom of entrepreneurs to make use of technology to create wealth.

But this kind of freedom will be different from all other freedoms that people have previously enjoyed. Political freedom has heretofore meant the freedom to pursue those ends that our natures had established for us. Those ends are not rigidly determined; human nature is very plastic, and we have an enormous range of choices conformable with that nature. But it is not infinitely malleable, and the elements that remain constant—particularly our species-typical gamut of emotional responses—constitute a safe harbor that allows us to connect, potentially, with all other human beings.

It may be that we are somehow destined to take up this new kind of freedom, or that the next stage of evolution is one in which, as some have suggested, we will deliberately take charge of our own biological makeup rather than leaving it to the blind forces of natural selection. But if we do, we should do it with eyes open. Many assume that the posthuman world will look pretty much like our own—free, equal, prosperous, caring, compassionate—only with better health care, longer lives, and perhaps more intelligence than today.

But the posthuman world could be one that is far more hierarchical and competitive than the one that currently exists, and full of social conflict as a result. It could be one in which any notion of "shared humanity" is lost, because we have mixed human genes with those of so many other species that we no longer have a clear idea of what a human being is. It could be one in which the median person is living well into his or her second century, sitting in a nursing home hoping for an unattainable death. Or it could be the kind of soft tyranny envisioned in *Brave New World*, in which everyone is healthy and happy but has forgotten the meaning of hope, fear, or struggle.

We do not have to accept any of these future worlds under a false banner of liberty, be it that of unlimited reproductive rights or of unfettered scientific inquiry. We do not have to regard ourselves as slaves to inevitable technological progress when that progress does not serve human ends. True freedom means the freedom of political communities to protect the values they hold most dear, and it is that freedom that we need to exercise with regard to the biotechnology revolution today.

2.5.2 PREVENTING A BRAVE NEW WORLD

LEON R. KASS

In this selection, Leon Kass, Addie Clark Harding Professor at the Committee on Social Thought at the University of Chicago and member of the President's Commission on Bioethics, argues that modern medical science is poised to cross an ethical boundary that will have momentous consequences for the future of humanity. Harking back to the dystopian vision of Aldous Huxley's classic *Brave New World* (1932), Kass argues that "the technological imperative, liberal democratic society, compassionate humanitarianism, moral pluralism, and free markets" are leading us down a path that places us at risk of losing our humanity. Kass argues that we should enact a worldwide ban on human cloning as a means of deterring "renegade scientists" from engaging in the practice. His proposed ban would apply to both reproductive and therapeutic cloning of human embryos because he believes that banning only reproductive cloning would prove impossible to enforce. Kass believes that we have the power to exercise control over the technological project but can only do so if we muster the political will to just say "no" to human cloning.

FOCUS QUESTIONS

1. What are the three main factors that Kass identifies as limiting our ability to control the onward march of the biomedical project? What other factors contribute to our inability to avoid the dangers Kass is concerned about?
2. What reasons does Kass give for "drawing the line" at human reproductive cloning by prohibiting the practice? Are these reasons convincing? Do the same arguments apply with equal force to therapeutic cloning? Explain.
3. Why does Kass believe that employing cloning or genetic engineering techniques to produce human children is "profoundly dehumanizing, no matter how good the product?" Explain.
4. Is there a moral difference between applying genetic engineering techniques to humans and to plants and nonhuman animals? Explain.

KEYWORDS

dehumanization, genetic determinism, infertility, reproductive cloning, reproductive freedom, somatic cell nuclear transfer

Source: From "Preventing a Brave New World," by Leon R. Kass, M.D., *The New Republic*, May 21, 2001. Copyright © 2001 by Leon R. Kass. Reprinted by permission of the author.

I.

The urgency of the great political struggles of the twentieth century, successfully waged against totalitarianisms first right and then left, seems to have blinded many people to a deeper and ultimately darker truth about the present age: all contemporary societies are traveling briskly in the same utopian direction. All are wedded to the modern technological project; all march eagerly to the drums of progress and fly proudly the banner of modern science; all sing loudly the Baconian anthem, "Conquer nature, relieve man's estate." Leading the triumphal procession is modern medicine, which is daily becoming ever more powerful in its battle against disease, decay, and death, thanks especially to astonishing achievements in biomedical science and technology—achievements for which we must surely be grateful.

Yet contemplating present and projected advances in genetic and reproductive technologies, in neuroscience and psychopharmacology, and in the development of artificial organs and computer-chip implants for human brains, we now clearly recognize new uses for biotechnical power that soar beyond the traditional medical goals of healing disease and relieving suffering. Human nature itself lies on the operating table, ready for alteration, for eugenic and psychic "enhancement," for wholesale re-design. In leading laboratories, academic and industrial, new creators are confidently amassing their powers and quietly honing their skills, while on the street their evangelists are zealously prophesying a post-human future. For anyone who cares about preserving our humanity, the time has come to pay attention.

Some transforming powers are already here. The Pill. In vitro fertilization. Bottled embryos. Surrogate wombs. Cloning. Genetic screening. Genetic manipulation. Organ harvesting. Mechanical spare parts. Chimeras. Brain implants. Ritalin for the young, Viagra for the old, Prozac for everyone. And, to leave this vale of tears, a little extra morphine accompanied by Muzak.

Years ago Aldous Huxley saw it coming. In his charming but disturbing novel, *Brave New World* (it appeared in 1932 and is more powerful on each re-reading), he made its meaning strikingly visible for all to see. Unlike other frightening futuristic novels of the past century, such as Orwell's already dated *Nineteen Eighty-Four*, Huxley shows us a dystopia that goes with, rather than against, the human grain. Indeed, it is animated by our own most humane and progressive aspirations. Following those aspirations to their ultimate realization, Huxley enables us to recognize those less obvious but often more pernicious evils that are inextricably linked to the successful attainment of partial goods.

Huxley depicts human life seven centuries hence, living under the gentle hand of humanitarianism rendered fully competent by genetic manipulation, psychoactive drugs, hypnopedia, and high-tech amusements. At long last, mankind has succeeded in eliminating disease, aggression, war, anxiety, suffering, guilt, envy, and grief. But this victory comes at the heavy price of homogenization, mediocrity, trivial pursuits, shallow attachments, debased tastes, spurious contentment, and souls without loves or longings. The Brave New World has achieved prosperity, community, stability, and nigh-universal contentment, only to be peopled by creatures of human shape but stunted humanity. They consume, fornicate, take "soma," enjoy "centrifugal bumble-puppy," and operate the machinery that makes it all possible. They do not read, write, think, love, or govern themselves. Art and science, virtue and religion, family and friendship are all passé. What matters most is bodily health and immediate gratification: "Never put off till tomorrow the fun you can have today." Brave New Man is so dehumanized that he does not even recognize what has been lost.

Huxley's novel, of course, is science fiction. Prozac is not yet Huxley's "soma"; cloning by nuclear transfer or splitting embryos is not exactly "Bokanovskification"; MTV and virtual-reality parlors are not quite the "feelies"; and our current safe and consequenceless sexual practices are not universally as loveless or as empty as those in the novel. But the kinships are disquieting, all

the more so since our technologies of bio-psycho-engineering are still in their infancy, and in ways that make all too clear what they might look like in their full maturity. Moreover, the cultural changes that technology has already wrought among us should make us even more worried than Huxley would have us be.

In Huxley's novel, everything proceeds under the direction of an omnipotent—albeit benevolent—world state. Yet the dehumanization that he portrays does not really require despotism or external control. To the contrary, precisely because the society of the future will deliver exactly what we most want—health, safety, comfort, plenty, pleasure, peace of mind and length of days—we can reach the same humanly debased condition solely on the basis of free human choice. No need for World Controllers. Just give us the technological imperative, liberal democratic society, compassionate humanitarianism, moral pluralism, and free markets, and we can take ourselves to a Brave New World all by ourselves—and without even deliberately deciding to go. In case you had not noticed, the train has already left the station and is gathering speed, but no one seems to be in charge.

Some among us are delighted, of course, by this state of affairs: some scientists and biotechnologists, their entrepreneurial backers, and a cheering claque of sci-fi enthusiasts, futurologists, and libertarians. There are dreams to be realized, powers to be exercised, honors to be won, and money—big money—to be made. But many of us are worried, and not, as the proponents of the revolution self-servingly claim, because we are either ignorant of science or afraid of the unknown. To the contrary, we can see all too clearly where the train is headed, and we do not like the destination. We can distinguish cleverness about means from wisdom about ends, and we are loath to entrust the future of the race to those who cannot tell the difference. No friend of humanity cheers for a post-human future.

Yet for all our disquiet, we have until now done nothing to prevent it. We hide our heads in the sand because we enjoy the blessings that medicine keeps supplying, or we rationalize our inaction by declaring that human engineering is inevitable and we can do nothing about it. In either case, we are complicit in preparing for our own degradation, in some respects more to blame than the bio-zealots who, however misguided, are putting their money where their mouth is. Denial and despair, unattractive outlooks in any situation, become morally reprehensible when circumstances summon us to keep the world safe for human flourishing.

Our immediate ancestors, taking up the challenge of their time, rose to the occasion and rescued the human future from the cruel dehumanizations of Nazi and Soviet tyranny. It is our more difficult task to find ways to preserve it from the soft dehumanizations of well-meaning but hubristic biotechnical "re-creationism"—and to do it without undermining biomedical science or rejecting its genuine contributions to human welfare.

Truth be told, it will not be easy for us to do so, and we know it. But rising to the challenge requires recognizing the difficulties. For there are indeed many features of modern life that will conspire to frustrate efforts aimed at the human control of the biomedical project. First, we Americans believe in technological automatism: where we do not foolishly believe that all innovation is progress, we fatalistically believe that it is inevitable ("If it can be done, it will be done, like it or not"). Second, we believe in freedom: the freedom of scientists to inquire, the freedom of technologists to develop, the freedom of entrepreneurs to invest and to profit, the freedom of private citizens to make use of existing technologies to satisfy any and all personal desires, including the desire to reproduce by whatever means. Third, the biomedical enterprise occupies the moral high ground of compassionate humanitarianism, upholding the supreme values of modern life—cure disease, prolong life, relieve suffering—in competition with which other moral goods rarely stand a chance. ("What the public wants is not to be sick," says James Watson, "and if we help them not to be sick, they'll be on our side.")

There are still other obstacles. Our cultural pluralism and easygoing relativism make it difficult to reach consensus on what we should embrace and what we should oppose; and moral objections

to this or that biomedical practice are often facilely dismissed as religious or sectarian. Many people are unwilling to pronounce judgments about what is good or bad, right and wrong, even in matters of great importance, even for themselves—never mind for others or for society as a whole. It does not help that the biomedical project is now deeply entangled with commerce: there are increasingly powerful economic interests in favor of going full steam ahead, and no economic interests in favor of going slow. Since we live in a democracy, moreover, we face political difficulties in gaining a consensus to direct our future, and we have almost no political experience in trying to curtail the development of any new biomedical technology. Finally, and perhaps most troubling, our views of the meaning of our humanity have been so transformed by the scientific-technological approach to the world that we are in danger of forgetting what we have to lose, humanly speaking.

But though the difficulties are real, our situation is far from hopeless. Regarding each of the aforementioned impediments, there is another side to the story. Though we love our gadgets and believe in progress, we have lost our innocence regarding technology. The environmental movement especially has alerted us to the unintended damage caused by unregulated technological advance, and has taught us how certain dangerous practices can be curbed. Though we favor freedom of inquiry, we recognize that experiments are deeds and not speeches, and we prohibit experimentation on human subjects without their consent, even when cures from disease might be had by unfettered research; and we limit so-called reproductive freedom by proscribing incest, polygamy, and the buying and selling of babies.

Although we esteem medical progress, biomedical institutions have ethics committees that judge research proposals on moral grounds, and, when necessary, uphold the primacy of human freedom and human dignity even over scientific discovery. Our moral pluralism notwithstanding, national commissions and review bodies have sometimes reached moral consensus to recommend against this or that technological application. On the economic front, the patenting of genes and life forms and the rapid rise of genomic commerce have elicited strong concerns and criticisms, leading even former enthusiasts of the new biology to recoil from the impending commodification of human life. Though we lack political institutions experienced in setting limits on biomedical innovation, federal agencies years ago rejected the development of the plutonium-powered artificial heart, and we have nationally prohibited commercial traffic in organs for transplantation, even though a market would increase the needed supply. In recent years, several American states and many foreign countries have successfully taken political action, making certain practices illegal and placing others under moratoriums (the creation of human embryos solely for research; human germ-line genetic alteration). Most importantly, the majority of Americans are not yet so degraded or so cynical as to fail to be revolted by the society depicted in Huxley's novel. Though the obstacles to effective action are significant, they offer no excuse for resignation. Besides, it would be disgraceful to concede defeat even before we enter the fray.

Not the least of our difficulties in trying to exercise control over where biology is taking us is the fact that we do not get to decide, once and for all, for or against the destination of a post-human world. The scientific discoveries and the technical powers that will take us there come to us piecemeal, one at a time and seemingly independent from one another, each often attractively introduced as a measure that will "help [us] not to be sick." But sometimes we come to a clear fork in the road where decision is possible, and where we know that our decision will make a world of difference—indeed, it will make a permanently different world. Fortunately, we stand now at the point of such a momentous decision. Events have conspired to provide us with a perfect opportunity to seize the initiative and to gain some control of the bio technical project. I refer to the prospect of human cloning, a practice absolutely central to Huxley's fictional world. Indeed, creating and manipulating life in the laboratory is the gateway to a Brave New World,

"To clone or not to clone a human being" is no longer a fanciful question. Success in cloning sheep, and also cows, mice, pigs, and goats, makes it perfectly clear that a fateful decision is now at hand: whether we should welcome or even tolerate the cloning of human beings. If recent newspaper reports are to be believed, reputable scientists and physicians have announced their intention to produce the first human clone in the coming year. Their efforts may already be under way.

The media, gawking and titillating as is their wont, have been softening us up for this possibility by turning the bizarre into the familiar. In the four years since the birth of Dolly the cloned sheep, the tone of discussing the prospect of human cloning has gone from "Yuck" to "Oh?" to "Gee whiz" to "Why not?" The sentimentalizers, aided by leading bioethicists, have downplayed talk about eugenically cloning the beautiful and the brawny or the best and the brightest. They have taken instead to defending clonal reproduction for humanitarian or compassionate reasons: to treat infertility in people who are said to "have no other choice," to avoid the risk of severe genetic disease, to "replace" a child who has died. For the sake of these rare benefits, they would have us countenance the entire practice of human cloning, the consequences be damned.

But we dare not be complacent about what is at issue, for the stakes are very high. Human cloning, though partly continuous with previous reproductive technologies, is also something radically new in itself and in its easily foreseeable consequences—especially when coupled with powers for genetic "enhancement" and germline genetic modification that may soon become available, owing to the recently completed Human Genome Project. I exaggerate somewhat, but in the direction of the truth: we are compelled to decide nothing less than whether human procreation is going to remain human, whether children are going to be made to order rather than begotten, and whether we wish to say yes in principle to the road that leads to the dehumanized hell of *Brave New World*.

...

For we have here a golden opportunity to exercise some control over where biology is taking us. The technology of cloning is discrete and well defined, and it requires considerable technical know-how and dexterity; we can therefore know by name many of the likely practitioners. The public demand for cloning is extremely low, and most people are decidedly against it. Nothing scientifically or medically important would be lost by banning clonal reproduction; alternative and non-objectionable means are available to obtain some of the most important medical benefits claimed for (nonreproductive) human cloning.

The commercial interests in human cloning are, for now, quite limited; and the nations of the world are actively seeking to prevent it. Now may be as good a chance as we will ever have to get our hands on the wheel of the runaway train now headed for a post-human world and to steer it toward a more dignified human future.

II.

What is cloning? Cloning, or asexual reproduction, is the production of individuals who are genetically identical to an already existing individual. The procedure's name is fancy—"somatic cell nuclear transfer"—but its concept is simple. Take a mature but unfertilized egg; remove or deactivate its nucleus; introduce a nucleus obtained from a specialized (somatic) cell of an adult organism. Once the egg begins to divide, transfer the little embryo to a woman's uterus to initiate a pregnancy. Since almost all the hereditary material of a cell is contained within its nucleus, the re-nucleated egg and the individual into which it develops are genetically identical to the organism that was the source of the transferred nucleus.

An unlimited number of genetically identical individuals—the group, as well as each of its members, is called "a clone"—could be produced by nuclear transfer. In principle, any person, male or female, newborn or adult, could be cloned, and in any quantity; and because stored cells can outlive their sources, one may even clone the dead. Since cloning requires no personal involvement on the part of the person whose genetic

material is used, it could easily be used to reproduce living or deceased persons without their consent—a threat to reproductive freedom that has received relatively little attention.

Some possible misconceptions need to be avoided. Cloning is not Xeroxing: the clone of Bill Clinton, though his genetic double, would enter the world hairless, toothless, and peeing in his diapers, like any other human infant. But neither is cloning just like natural twinning: the cloned twin will be identical to an older, existing adult; and it will arise not by chance but by deliberate design; and its entire genetic makeup will be preselected by its parents and/or scientists.

Moreover, the success rate of cloning, at least at first, will probably not be very high: the Scots transferred two hundred seventy-seven adult nuclei into sheep eggs, implanted twenty-nine clonal embryos, and achieved the birth of only one live lamb clone.

For this reason, among others, it is unlikely that, at least for now, the practice would be very popular; and there is little immediate worry of mass-scale production of multicopies. Still, for the tens of thousands of people who sustain more than three hundred assisted-reproduction clinics in the United States and already avail themselves of in vitro fertilization and other techniques, cloning would be an option with virtually no added fuss. Panos Zavos, the Kentucky reproduction specialist who has announced his plans to clone a child, claims that he has already received thousands of e-mailed requests from people eager to clone, despite the known risks of failure and damaged offspring. Should commercial interests develop in "nucleus-banking," as they have in sperm-banking and egg-harvesting; should famous athletes or other celebrities decide to market their DNA the way they now market their autographs and nearly everything else; should techniques of embryo and germline genetic testing and manipulation arrive as anticipated, increasing the use of laboratory assistance in order to obtain "better" babies—should all this come to pass, cloning, if it is permitted, could become more than a marginal practice simply on the basis of free reproductive choice.

What are we to think about this prospect? Nothing good. Indeed, most people are repelled by nearly all aspects of human cloning: the possibility of mass production of human beings, with large clones of look-alikes, compromised in their individuality; the idea of father-son or mother-daughter "twins"; the bizarre prospect of a woman bearing and rearing a genetic copy of herself, her spouse, or even her deceased father or mother; the grotesqueness of conceiving a child as an exact "replacement" for another who has died; the utilitarian creation of embryonic duplicates of oneself, to be frozen away or created when needed to provide homologous tissues or organs for transplantation; the narcissism of those who would clone themselves, and the arrogance of others who think they know who deserves to be cloned; the Frankensteinian hubris to create a human life and increasingly to control its destiny; men playing at being God. Almost no one finds any of the suggested reasons for human cloning compelling, and almost everyone anticipates its possible misuses and abuses. And the popular belief that human cloning cannot be prevented makes the prospect all the more revolting.

Revulsion is not an argument; and some of yesterday's repugnances are today calmly accepted—not always for the better. In some crucial cases, however, repugnance is the emotional expression of deep wisdom, beyond reason's power completely to articulate it. Can anyone really give an argument fully adequate to the horror that is father-daughter incest (even with consent), or bestiality, or the mutilation of a corpse, or the eating of human flesh, or the rape or murder of another human being? Would anybody's failure to give full rational justification for his revulsion at those practices make that revulsion ethically suspect?

I suggest that our repugnance at human cloning belongs in this category. We are repelled by the prospect of cloning human beings not because of the strangeness or the novelty of the undertaking, but because we intuit and we feel, immediately and without argument, the violation of things that we rightfully hold dear. We sense that cloning represents a profound defilement of

our given nature as procreative beings, and of the social relations built on this natural ground. We also sense that cloning is a radical form of child abuse. In this age in which everything is held to be permissible so long as it is freely done, and in which our bodies are regarded as mere instruments of our autonomous rational will, repugnance may be the only voice left that speaks up to defend the central core of our humanity. Shallow are the souls that have forgotten how to shudder.

III.

Yet repugnance need not stand naked before the bar of reason. The wisdom of our horror at human cloning can be at least partially articulated, even if this is finally one of those instances about which the heart has its reasons that reason cannot entirely know. I offer four objections to human cloning: that it constitutes unethical experimentation; that it threatens identity and individuality; that it turns procreation into manufacture (especially when understood as the harbinger of manipulations to come); and that it means despotism over children and perversion of parenthood. Please note: I speak only about so-called reproductive cloning, not about the creation of cloned embryos for research. The objections that may be raised against creating (or using) embryos for research are entirely independent of whether the research embryos are produced by cloning. What is radically distinct and radically new is reproductive cloning.

Any attempt to clone a human being would constitute an unethical experiment upon the resulting child-to-be. In all the animal experiments, fewer than two to three percent of all cloning attempts succeeded. Not only are there fetal deaths and stillborn infants, but many of the so-called "successes" are in fact failures. As has only recently become clear, there is a very high incidence of major disabilities and deformities in cloned animals that attain live birth. Cloned cows often have heart and lung problems; cloned mice later develop pathological obesity; other live-born cloned animals fail to reach normal developmental milestones.

The problem, scientists suggest, may lie in the fact that an egg with a new somatic nucleus must re-program itself in a matter of minutes or hours (whereas the nucleus of an unaltered egg has been prepared over months and years). There is thus a greatly increased likelihood of error in translating the genetic instructions, leading to developmental defects some of which will show themselves only much later. (Note also that these induced abnormalities may also affect the stem cells that scientists hope to harvest from cloned embryos. Lousy embryos, lousy stem cells.) Nearly all scientists now agree that attempts to clone human beings carry massive risks of producing unhealthy, abnormal, and malformed children. What are we to do with them? Shall we just discard the ones that fall short of expectations? Considered opinion is today nearly unanimous, even among scientists: attempts at human cloning are irresponsible and unethical. We cannot ethically even get to know whether or not human cloning is feasible.

If it were successful, cloning would create serious issues of identity and individuality. The clone may experience concerns about his distinctive identity not only because he will be, in genotype and in appearance, identical to another human being, but because he may also be twin to the person who is his "father" or his "mother"—if one can still call them that. Unaccountably, people treat as innocent the homey case of intra-familial cloning—the cloning of husband or wife (or single mother). They forget about the unique dangers of mixing the twin relation with the parent-child relation. (For this situation, the relation of contemporaneous twins is no precedent; yet even this less problematic situation teaches us how difficult it is to wrest independence from the being for whom one has the most powerful affinity.) Virtually no parent is going to be able to treat a clone of himself or herself as one treats a child generated by the lottery of sex. What will happen when the adolescent clone of Mommy becomes the spitting image of the woman with whom Daddy once fell in love? In case of divorce, will Mommy still love the clone of Daddy, even

though she can no longer stand the sight of Daddy himself?

Most people think about cloning from the point of view of adults choosing to clone. Almost nobody thinks about what it would be like to be the cloned child. Surely his or her new life would constantly be scrutinized in relation to that of the older version. Even in the absence of unusual parental expectations for the clone—say, to live the same life, only without its errors—the child is likely to be ever a curiosity, ever a potential source of déjà vu. Unlike "normal" identical twins, a cloned individual—copied from whomever—will be saddled with a genotype that has already lived. He will not be fully a surprise to the world: people are likely always to compare his doings in life with those of his alter ego, especially if he is a clone of someone gifted or famous. True, his nurture and his circumstance will be different; genotype is not exactly destiny. But one must also expect parental efforts to shape this new life after the original—or at least to view the child with the original version always firmly in mind. For why else did they clone from the star basketball player, the mathematician, or the beauty queen—or even dear old Dad—in the first place?

Human cloning would also represent a giant step toward the transformation of begetting into making, of procreation into manufacture (literally, "handmade"), a process that has already begun with in vitro fertilization and genetic testing of embryos. With cloning, not only is the process in hand, but the total genetic blueprint of the cloned individual is selected and determined by the human artisans. To be sure, subsequent development is still according to natural processes; and the resulting children will be recognizably human. But we would be taking a major step into making man himself simply another one of the man-made things.

How does begetting differ from making? In natural procreation, human beings come together to give existence to another being that is formed exactly as we were, by what we are—living, hence perishable, hence aspiringly erotic, hence procreative human beings. But in clonal reproduction, and in the more advanced forms of manufacture to which it will lead, we give existence to a being not by what we are but by what we intend and design.

Let me be clear. The problem is not the mere intervention of technique, and the point is not that "nature knows best." The problem is that any child whose being, character, and capacities exist owing to human design does not stand on the same plane as its makers. As with any product of our making, no matter how excellent, the artificer stands above it, not as an equal but as a superior, transcending it by his will and creative prowess. In human cloning, scientists and prospective "parents" adopt a technocratic attitude toward human children: human children become their artifacts. Such an arrangement is profoundly dehumanizing, no matter how good the product.

Procreation dehumanized into manufacture is further degraded by commodification, a virtually inescapable result of allowing baby-making to proceed under the banner of commerce. Genetic and reproductive biotechnology companies are already growth industries, but they will soon go into commercial orbit now that the Human Genome Project has been completed. "Human eggs for sale" is already a big business, masquerading under the pretense of "donation." Newspaper advertisements on elite college campuses offer up to $50,000 for an egg "donor" tall enough to play women's basketball and with SAT scores high enough for admission to Stanford; and to nobody's surprise, at such prices there are many young coeds eager to help shoppers obtain the finest babies money can buy. (The egg and womb-renting entrepreneurs shamelessly proceed on the ancient, disgusting, misogynist premise that most women will give you access to their bodies, if the price is right.) Even before the capacity for human cloning is perfected, established companies will have invested in the harvesting of eggs from ovaries obtained at autopsy or through ovarian surgery, practiced embryonic genetic alteration, and initiated the stockpiling of prospective donor tissues. Through the rental of surrogate-womb services, and through the buying and selling of tissues and embryos priced according to the merit of the donor, the commodification of nascent human life will be unstoppable.

Finally, the practice of human cloning by nuclear transfer—like other anticipated forms of genetically engineering the next generation—would enshrine and aggravate a profound misunderstanding of the meaning of having children and of the parent-child relationship. When a couple normally chooses to procreate, the partners are saying yes to the emergence of new life in its novelty—are saying yes not only to having a child, but also to having whatever child this child turns out to be. In accepting our finitude, in opening ourselves to our replacement, we tacitly confess the limits of our control.

Embracing the future by procreating means precisely that we are relinquishing our grip in the very activity of taking up our own share in what we hope will be the immortality of human life and the human species. This means that our children are not our children: they are not our property, they are not our possessions. Neither are they supposed to live our lives for us, or to live anyone's life but their own. Their genetic distinctiveness and independence are the natural foreshadowing of the deep truth that they have their own, never-before-enacted life to live. Though sprung from a past, they take an uncharted course into the future.

Much mischief is already done by parents who try to live vicariously through their children. Children are sometimes compelled to fulfill the broken dreams of unhappy parents. But whereas most parents normally have hopes for their children, cloning parents will have expectations. In cloning, such overbearing parents will have taken at the start a decisive step that contradicts the entire meaning of the open and forward-looking nature of parent-child relations. The child is given a genotype that has already lived, with full expectation that this blueprint of a past life ought to be controlling the life that is to come. A wanted child now means a child who exists precisely to fulfill parental wants. Like all the more precise eugenic manipulations that will follow in its wake, cloning is thus inherently despotic, for it seeks to make one's children after one's own image (or an image of one's choosing) and their future according to one's will.

Is this hyperbolic? Consider concretely the new realities of responsibility and guilt in the households of the cloned. No longer only the sins of the parents, but also the genetic choices of the parents, will be visited on the children—and beyond the third and fourth generation; and everyone will know who is responsible. No parent will be able to blame nature or the lottery of sex for an unhappy adolescent's big nose, dull wit, musical ineptitude, nervous disposition, or anything else that he hates about himself. Fairly or not, children will hold their cloners responsible for everything, for nature as well as for nurture. And parents, especially the better ones, will be limitlessly liable to guilt. Only the truly despotic souls will sleep the sleep of the innocent.

IV.

The defenders of cloning are not wittingly friends of despotism. Quite the contrary. Deaf to most other considerations, they regard themselves mainly as friends of freedom: the freedom of individuals to reproduce, the freedom of scientists and inventors to discover and to devise and to foster "progress" in genetic knowledge and technique, the freedom of entrepreneurs to profit in the market. They want large-scale cloning only for animals, but they wish to preserve cloning as a human option for exercising our "right to reproduce"—our right to have children, and children with "desirable genes." As some point out, under our "right to reproduce" we already practice early forms of unnatural, artificial, and extramarital reproduction, and we already practice early forms of eugenic choice. For that reason, they argue, cloning is no big deal.

We have here a perfect example of the logic of the slippery slope. The principle of reproductive freedom currently enunciated by the proponents of cloning logically embraces the ethical acceptability of sliding all the way down: to producing children wholly in the laboratory from sperm to term (should it become feasible), and to producing children whose entire genetic makeup will be the product of parental eugenic planning and choice. If reproductive freedom means the right to have a child of one's own choosing by whatever

means, then reproductive freedom knows and accepts no limits.

Proponents want us to believe that there are legitimate uses of cloning that can be distinguished from illegitimate uses, but by their own principles no such limits can be found. (Nor could any such limits be enforced in practice: once cloning is permitted, no one ever need discover whom one is cloning and why.) Reproductive freedom, as they understand it, is governed solely by the subjective wishes of the parents-to-be. The sentimentally appealing case of the childless married couple is, on these grounds, indistinguishable from the case of an individual (married or not) who would like to clone someone famous or talented, living or dead. And the principle here endorsed justifies not only cloning but also all future artificial attempts to create (manufacture) "better" or "perfect" babies.

The "perfect baby," of course, is the project not of the infertility doctors, but of the eugenic scientists and their supporters, who, for the time being, are content to hide behind the skirts of the partisans of reproductive freedom and compassion for the infertile. For them, the paramount right is not the so-called right to reproduce, it is what the biologist Bentley Glass called, a quarter of a century ago, "the right of every child to be born with a sound physical and mental constitution, based on a sound genotype … the inalienable right to a sound heritage." But to secure this right, and to achieve the requisite quality control over new human life, human conception and gestation will need to be brought fully into the bright light of the laboratory, beneath which the child-to-be can be fertilized, nourished, pruned, weeded, watched, inspected, prodded, pinched, cajoled, injected, tested, rated, graded, approved, stamped, wrapped, sealed, and delivered. There is no other way to produce the perfect baby.

If you think that such scenarios require outside coercion or governmental tyranny, you are mistaken. Once it becomes possible, with the aid of human genomics, to produce or to select for what some regard as "better babies"—smarter, prettier, healthier, more athletic—parents will leap at the opportunity to "improve" their offspring. Indeed, not to do so will be socially regarded as a form of child neglect. Those who would ordinarily be opposed to such tinkering will be under enormous pressure to compete on behalf of their as yet unborn children—just as some now plan almost from their children's birth how to get them into Harvard. Never mind that, lacking a standard of "good" or "better," no one can really know whether any such changes will truly be improvements.

Proponents of cloning urge us to forget about the science-fiction scenarios of laboratory manufacture or multiple-copy clones, and to focus only on the sympathetic cases of infertile couples exercising their reproductive rights. But why, if the single cases are so innocent, should multiplying their performance be so off-putting? (Similarly, why do others object to people's making money from that practice if the practice itself is perfectly acceptable?) The so-called science-fiction cases—say, Brave New World—make vivid the meaning of what looks to us, mistakenly, to be benign. They reveal that what looks like compassionate humanitarianism is, in the end, crushing dehumanization.

V.

Whether or not they share my reasons, most people, I think, share my conclusion: that human cloning is unethical in itself and dangerous in its likely consequences, which include the precedent that it will establish for designing our children. Some reach this conclusion for their own good reasons, different from my own: concerns about distributive justice in access to eugenic cloning; worries about the genetic effects of asexual "inbreeding"; aversion to the implicit premise of genetic determinism; objections to the embryonic and fetal wastage that must necessarily accompany the efforts; religious opposition to "man playing God." But never mind why: the overwhelming majority of our fellow Americans remain firmly opposed to cloning human beings.

For us, then, the real questions are: What should we do about it? How can we best succeed? These questions should concern everyone eager to secure deliberate human control over the powers that could re-design our humanity, even if cloning

is not the issue over which they would choose to make their stand. And the answer to the first question seems pretty plain. What we should do is work to prevent human cloning by making it illegal.

We should aim for a global legal ban, if possible, and for a unilateral national ban at a minimum—and soon, before the fact is upon us. To be sure, legal bans can be violated; but we certainly curtail much mischief by outlawing incest, voluntary servitude, and the buying and selling of organs and babies. To be sure, renegade scientists may secretly undertake to violate such a law, but we can deter them by both criminal sanctions and monetary penalties, as well as by removing any incentive they have to proudly claim credit for their technological bravado.

Such a ban on clonal baby-making will not harm the progress of basic genetic science and technology. On the contrary, it will reassure the public that scientists are happy to proceed without violating the deep ethical norms and intuitions of the human community. It will also protect honorable scientists from a public backlash against the brazen misconduct of the rogues. As many scientists have publicly confessed, free and worthy science probably has much more to fear from a strong public reaction to a cloning fiasco than it does from a cloning ban, provided that the ban is judiciously crafted and vigorously enforced against those who would violate it....

... I now believe that what we need is an all-out ban on human cloning, including the creation of embryonic clones. I am convinced that all halfway measures will prove to be morally, legally, and strategically flawed, and—most important—that they will not be effective in obtaining the desired result. Anyone truly serious about preventing human reproductive cloning must seek to stop the process from the beginning. Our changed circumstances, and the now evident defects of the less restrictive alternatives, make an all-out ban by far the most attractive and effective option.

Here's why. Creating cloned human children ("reproductive cloning") necessarily begins by producing cloned human embryos. Preventing the latter would prevent the former, and prudence alone might counsel building such a "fence around the law." Yet some scientists favor embryo cloning as a way of obtaining embryos for research or as sources of cells and tissues for the possible benefit of others. (This practice they misleadingly call "therapeutic cloning" rather than the more accurate "cloning for research" or "experimental cloning," so as to obscure the fact that the clone will be "treated" only to exploitation and destruction, and that any potential future beneficiaries and any future "therapies" are at this point purely hypothetical.)

The prospect of creating new human life solely to be exploited in this way has been condemned on moral grounds by many people—including *The Washington Post*, President Clinton, and many other supporters of a woman's right to abortion—as displaying a profound disrespect for life. Even those who are willing to scavenge so-called "spare embryos"—those products of in vitro fertilization made in excess of people's reproductive needs, and otherwise likely to be discarded—draw back from creating human embryos explicitly and solely for research purposes.

They reject outright what they regard as the exploitation and the instrumentalization of nascent human life. In addition, others who are agnostic about the moral status of the embryo see the wisdom of not needlessly offending the sensibilities of their fellow citizens who are opposed to such practices.

But even setting aside these obvious moral first impressions, a few moments of reflection show why an anti-cloning law that permitted the cloning of embryos but criminalized their transfer to produce a child would be a moral blunder. This would be a law that was not merely permissively "pro-choice" but emphatically and prescriptively "anti-life." While permitting the creation of an embryonic life, it would make it a federal offense to try to keep it alive and bring it to birth. Whatever one thinks of the moral status or the ontological status of the human embryo, moral sense and practical wisdom recoil from having the government of the United States on record as requiring the destruction of nascent life and, what is worse, demanding the punishment of those who would act to preserve it by (feloniously!) giving it birth.

But the problem with the approach that targets only reproductive cloning (that is, the transfer of the embryo to a woman's uterus) is not only moral but also legal and strategic. A ban only on reproductive cloning would turn out to be unenforceable. Once cloned embryos were produced and available in laboratories and assisted-reproduction centers, it would be virtually impossible to control what was done with them. Biotechnical experiments take place in laboratories, hidden from public view, and, given the rise of high-stakes commerce in biotechnology, these experiments are concealed from the competition. Huge stockpiles of cloned human embryos could thus be produced and bought and sold without anyone knowing it. As we have seen with in vitro embryos created to treat infertility, embryos produced for one reason can be used for another reason: today "spare embryos" once created to begin a pregnancy are now used in research, and tomorrow clones created for research will be used to begin a pregnancy.

Assisted reproduction takes place within the privacy of the doctor-patient relationship, making outside scrutiny extremely difficult. Many infertility experts probably would obey the law, but others could and would defy it with impunity, their doings covered by the veil of secrecy that is the principle of medical confidentiality. Moreover, the transfer of embryos to begin a pregnancy is a simple procedure (especially compared with manufacturing the embryo in the first place), simple enough that its final steps could be self-administered by the woman, who would thus absolve the doctor of blame for having "caused" the illegal transfer. (I have in mind something analogous to Kevorkian's suicide machine, which was designed to enable the patient to push the plunger and the good "doctor" to evade criminal liability.)

Even should the deed become known, governmental attempts to enforce the reproductive ban would run into a swarm of moral and legal challenges, both to efforts aimed at preventing transfer to a woman and—even worse—to efforts seeking to prevent birth after transfer has occurred. A woman who wished to receive the embryo clone would no doubt seek a judicial restraining order, suing to have the law overturned in the name of a constitutionally protected interest in her own reproductive choice to clone. (The cloned child would be born before the legal proceedings were complete.) And should an "illicit clonal pregnancy" be discovered, no governmental agency would compel a woman to abort the clone, and there would be an understandable storm of protest should she be fined or jailed after she gives birth. Once the baby is born, there would even be sentimental opposition to punishing the doctor for violating the law—unless, of course, the clone turned out to be severely abnormal.

For all these reasons, the only practically effective and legally sound approach is to block human cloning at the start, at the production of the embryo clone. Such a ban can be rightly characterized not as interference with reproductive freedom, nor even as interference with scientific inquiry, but as an attempt to prevent the unhealthy, unsavory, and unwelcome manufacture of and traffic in human clones.

...

I appreciate that a federal legislative ban on human cloning is without American precedent, at least in matters technological. Perhaps such a ban will prove ineffective; perhaps it will eventually be shown to have been a mistake. (If so, it could later be reversed.) If enacted, however, it will have achieved one overwhelmingly important result, in addition to its contribution to thwarting cloning: it will place the burden of practical proof where it belongs. It will require the proponents to show very clearly what great social or medical good can be had only by the cloning of human beings. Surely it is only for such a compelling case, yet to be made or even imagined, that we should wish to risk this major departure—or any other major departure—in human procreation.

Americans have lived by and prospered under a rosy optimism about scientific and technological progress. The technological imperative has probably served us well, though we should admit that there is no accurate method for weighing benefits and harms. And even when we recognize the unwelcome outcomes of technological

advance, we remain confident in our ability to fix all the "bad" consequences—by regulation or by means of still newer and better technologies. Yet there is very good reason for shifting the American paradigm, at least regarding those technological interventions into the human body and mind that would surely effect fundamental (and likely irreversible) changes in human nature, basic human relationships, and what it means to be a human being. Here we should not be willing to risk everything in the naive hope that, should things go wrong, we can later set them right again.

Some have argued that cloning is almost certainly going to remain a marginal practice, and that we should therefore permit people to practice it. Such a view is shortsighted. Even if cloning is rarely undertaken, a society in which it is tolerated is no longer the same society—any more than is a society that permits (even small-scale) incest or cannibalism or slavery. A society that allows cloning, whether it knows it or not, has tacitly assented to the conversion of procreation into manufacture and to the treatment of children as purely the projects of our will. Willy-nilly, it has acquiesced in the eugenic re-design of future generations. The humanitarian superhighway to a Brave New World lies open before this society.

But the present danger posed by human cloning is, paradoxically, also a golden opportunity. In a truly unprecedented way, we can strike a blow for the human control of the technological project, for wisdom, for prudence, for human dignity. The prospect of human cloning, so repulsive to contemplate, is the occasion for deciding whether we shall be slaves of unregulated innovation, and ultimately its artifacts, or whether we shall remain free human beings who guide our powers toward the enhancement of human dignity. The humanity of the human future is now in our hands.

2.5.3 THE CASE AGAINST PERFECTION

MICHAEL J. SANDEL

In this reading, Michael Sandel, professor of political philosophy at Harvard University, develops an answer to some of the hardest ethical questions raised by genetic engineering, particularly the use of this technology for human genetic enhancement. It is one thing, he argues, to employ genetic technologies to prevent or cure diseases or repair injuries but quite another to use them to produce children who are genetically advantaged by parental choice. Many people have moral qualms about performance-enhancing drugs in sports, but it is difficult to isolate exactly the ethical basis of these intuitions. Sandel examines and dismisses several plausible explanations for such qualms but finds each of these arguments insufficient. Instead, he argues that what is problematic about drug or genetic enhancement is that it represents "a Promethean aspiration to remake nature, including human nature, to serve our purposes and satisfy our desires." This desire to remake ourselves coexists with a countervailing value in the "giftedness of life"—that is, the appreciation that "our talents and powers

Source: From "The Case Against Perfection" by Michael J. Sandel, *The Atlantic Monthly*, April 2004, pp. 51–62. Copyright © 2004 The Atlantic Media Co., as first published in The Atlantic Magazine.

are not wholly our own doing." The "ethic of giftedness" provides a moral basis for rejecting the vision of "free-market eugenics" that has been advocated by many bioethicists. Instead, Sandel argues for constraining our drive to mastery with a humility born of the recognition that our individual genetic endowments should remain "gifts" rather than become products of deliberate design.

FOCUS QUESTIONS

1. Why does Sandel think that the objection to human cloning from autonomy or the value of an "open future" is unconvincing? What about the argument from unfairness? How can these objections be answered?
2. What does Sandel mean by the "ethic of willfulness?" How is this ethic manifest, for instance, in the use of performance-enhancing drugs by athletes?
3. What three key features of the moral landscape would be undermined, in Sandel's view, by the use of genetic enhancement technologies? Do you agree that these values should be protected and preserved? Explain.
4. Compare Sandel's view to those of Leon Kass (Selection 2.5.2). Is there any way to reconcile or balance their respective positions? Explain.

KEYWORDS

chromosomal lottery, eugenics, gene therapy, genetic enhancement, human growth hormone, performance-enhancing drugs, sex selection

Breakthroughs in genetics present us with a promise and a predicament. The promise is that we may soon be able to treat and prevent a host of debilitating diseases. The predicament is that our newfound genetic knowledge may also enable us to manipulate our own nature—to enhance our muscles, memories, and moods; to choose the sex, height, and other genetic traits of our children; to make ourselves "better than well." When science moves faster than moral understanding, as it does today, men and women struggle to articulate their unease. In liberal societies they reach first for the language of autonomy, fairness, and individual rights. But this part of our moral vocabulary is ill equipped to address the hardest questions posed by genetic engineering. The genomic revolution has induced a kind of moral vertigo.

Consider cloning. The birth of Dolly the cloned sheep, in 1997, brought a torrent of concern about the prospect of cloned human beings. There are good medical reasons to worry. Most scientists agree that cloning is unsafe, likely to produce offspring with serious abnormalities. (Dolly recently died a premature death.) But suppose technology improved to the point where clones were at no greater risk than naturally conceived offspring. Would human cloning still be objectionable? Should our hesitation be moral as well as medical? What, exactly, is wrong with creating a child who is a genetic twin of one parent, or of an older sibling who has tragically died—or, for that matter, of an admired scientist, sports star, or celebrity?

Some say cloning is wrong because it violates the right to autonomy: by choosing a child's genetic makeup in advance, parents deny the child's right to an open future. A similar objection can be raised against any form of bioengineering that allows parents to select or reject genetic characteristics. According to this argument, genetic enhancements for musical talent, say, or athletic prowess, would point children toward particular choices, and so designer children would never be fully free.

At first glance the autonomy argument seems to capture what is troubling about human cloning and other forms of genetic engineering. It is not persuasive, for two reasons. First, it wrongly implies that absent a designing parent, children

are free to choose their characteristics for themselves. But none of us chooses his genetic inheritance. The alternative to a cloned or genetically enhanced child is not one whose future is unbound by particular talents but one at the mercy of the genetic lottery.

Second, even if a concern for autonomy explains some of our worries about made-to-order children, it cannot explain our moral hesitation about people who seek genetic remedies or enhancements for themselves. Gene therapy on somatic (that is, non-reproductive) cells, such as muscle cells and brain cells, repairs or replaces defective genes. The moral quandary arises when people use such therapy not to cure a disease but to reach beyond health, to enhance their physical or cognitive capacities, to lift themselves above the norm.

Like cosmetic surgery, genetic enhancement employs medical means for nonmedical ends—ends unrelated to curing or preventing disease or repairing injury. But unlike cosmetic surgery, genetic enhancement is more than skin-deep. If we are ambivalent about surgery or Botox injections for sagging chins and furrowed brows, we are all the more troubled by genetic engineering for stronger bodies, sharper memories, greater intelligence, and happier moods. The question is whether we are right to be troubled, and if so, on what grounds.

In order to grapple with the ethics of enhancement, we need to confront questions largely lost from view—questions about the moral status of nature, and about the proper stance of human beings toward the given world. Since these questions verge on theology, modern philosophers and political theorists tend to shrink from them. But our new powers of biotechnology make them unavoidable. To see why this is so, consider four examples already on the horizon: muscle enhancement, memory enhancement, growth-hormone treatment, and reproductive technologies that enable parents to choose the sex and some genetic traits of their children. In each case what began as an attempt to treat a disease or prevent a genetic disorder now beckons as an instrument of improvement and consumer choice.

MUSCLES

Everyone would welcome a gene therapy to alleviate muscular dystrophy and to reverse the debilitating muscle loss that comes with old age. But what if the same therapy were used to improve athletic performance? Researchers have developed a synthetic gene that, when injected into the muscle cells of mice, prevents and even reverses natural muscle deterioration. The gene not only repairs wasted or injured muscles but also strengthens healthy ones. This success bodes well for human applications. H. Lee Sweeney, of the University of Pennsylvania, who leads the research, hopes his discovery will cure the immobility that afflicts the elderly. But Sweeney's bulked-up mice have already attracted the attention of athletes seeking a competitive edge. Although the therapy is not yet approved for human use, the prospect of genetically enhanced weight lifters, home-run sluggers, linebackers, and sprinters is easy to imagine. The widespread use of steroids and other performance-improving drugs in professional sports suggests that many athletes will be eager to avail themselves of genetic enhancement.

Suppose for the sake of argument that muscle-enhancing gene therapy, unlike steroids, turned out to be safe—or at least no riskier than a rigorous weight-training regimen. Would there be a reason to ban its use in sports? There is something unsettling about the image of genetically altered athletes lifting SUVs or hitting 650-foot home runs or running a three-minute mile. But what, exactly, is troubling about it? Is it simply that we find such superhuman spectacles too bizarre to contemplate? Or does our unease point to something of ethical significance?

It might be argued that a genetically enhanced athlete, like a drug-enhanced athlete, would have an unfair advantage over his unenhanced competitors. But the fairness argument against enhancement has a fatal flaw: it has always been the case that some athletes are better endowed genetically than others, and yet we do not consider this to undermine the fairness of competitive sports. From the standpoint of fairness, enhanced genetic differences would be

no worse than natural ones, assuming they were safe and made available to all. If genetic enhancement in sports is morally objectionable, it must be for reasons other than fairness.

MEMORY

Genetic enhancement is possible for brains as well as brawn. In the mid-1990s scientists managed to manipulate a memory-linked gene in fruit flies, creating flies with photographic memories. More recently researchers have produced smart mice by inserting extra copies of a memory-related gene into mouse embryos. The altered mice learn more quickly and remember things longer than normal mice. The extra copies were programmed to remain active even in old age, and the improvement was passed on to offspring. Human memory is more complicated, but biotech companies, including Memory Pharmaceuticals, are in hot pursuit of memory-enhancing drugs, or "cognition enhancers," for human beings. The obvious market for such drugs consists of those who suffer from Alzheimer's and other serious memory disorders. The companies also have their sights on a bigger market: the 81 million Americans over fifty who are beginning to encounter the memory loss that comes naturally with age. A drug that reversed age-related memory loss would be a bonanza for the pharmaceutical industry: a Viagra for the brain. Such use would straddle the line between remedy and enhancement. Unlike a treatment for Alzheimer's, it would cure no disease; but insofar as it restored capacities a person once possessed, it would have a remedial aspect. It could also have purely nonmedical uses: for example, by a lawyer cramming to memorize facts for an upcoming trial, or by a business executive eager to learn Mandarin on the eve of his departure for Shanghai.

Some who worry about the ethics of cognitive enhancement point to the danger of creating two classes of human beings: those with access to enhancement technologies, and those who must make do with their natural capacities. And if the enhancements could be passed down the generations, the two classes might eventually become subspecies—the enhanced and the merely natural. But worry about access ignores the moral status of enhancement itself. Is the scenario troubling because the unenhanced poor would be denied the benefits of bioengineering, or because the enhanced affluent would somehow be dehumanized? As with muscles, so with memory: the fundamental question is not how to ensure equal access to enhancement but whether we should aspire to it in the first place.

HEIGHT

Pediatricians already struggle with the ethics of enhancement when confronted by parents who want to make their children taller. Since the 1980s human growth hormone has been approved for children with a hormone deficiency that makes them much shorter than average. But the treatment also increases the height of healthy children. Some parents of healthy children who are unhappy with their stature (typically boys) ask why it should make a difference whether a child is short because of a hormone deficiency or because his parents happen to be short. Whatever the cause, the social consequences are the same.

In the face of this argument some doctors began prescribing hormone treatments for children whose short stature was unrelated to any medical problem. By 1996 such "off-label" use accounted for 40 percent of human-growth-hormone prescriptions. Although it is legal to prescribe drugs for purposes not approved by the Food and Drug Administration, pharmaceutical companies cannot promote such use. Seeking to expand its market, Eli Lilly & Co. recently persuaded the FDA to approve its human growth hormone for healthy children whose projected adult height is in the bottom one percentile—under five feet three inches for boys and four feet eleven inches for girls. This concession raises a large question about the ethics of enhancement: If hormone treatments need not be limited to those with hormone deficiencies, why should they be available only to very short children? Why shouldn't all shorter-than-average children be able to seek treatment? And what about a child of average height who wants to be taller so that he can make the basketball team?

Some oppose height enhancement on the grounds that it is collectively self-defeating; as some become taller, others become shorter relative to the norm. Except in Lake Wobegon, not every child can be above average. As the unenhanced began to feel shorter, they, too, might seek treatment, leading to a hormonal arms race that left everyone worse off, especially those who couldn't afford to buy their way up from shortness.

But the arms-race objection is not decisive on its own. Like the fairness objection to bioengineered muscles and memory, it leaves unexamined the attitudes and dispositions that prompt the drive for enhancement. If we were bothered only by the injustice of adding shortness to the problems of the poor, we could remedy that unfairness by publicly subsidizing height enhancements. As for the relative height deprivation suffered by innocent bystanders, we could compensate them by taxing those who buy their way to greater height. The real question is whether we want to live in a society where parents feel compelled to spend a fortune to make perfectly healthy kids a few inches taller.

SEX SELECTION

Perhaps the most inevitable nonmedical use of bioengineering is sex selection. For centuries parents have been trying to choose the sex of their children. Today biotech succeeds where folk remedies failed.

One technique for sex selection arose with prenatal tests using amniocentesis and ultrasound. These medical technologies were developed to detect genetic abnormalities such as spina bifida and Down syndrome. But they can also reveal the sex of the fetus—allowing for the abortion of a fetus of an undesired sex. Even among those who favor abortion rights, few advocate abortion simply because the parents do not want a girl. Nevertheless, in traditional societies with a powerful cultural preference for boys, this practice has become widespread.

Sex selection need not involve abortion, however. For couples undergoing *in vitro* fertilization (IVF), it is possible to choose the sex of the child before the fertilized egg is implanted in the womb. One method makes use of preimplantation genetic diagnosis (PGD), a procedure developed to screen for genetic diseases. Several eggs are fertilized in a petri dish and grown to the eight-cell stage (about three days). At that point the embryos are tested to determine their sex. Those of the desired sex are implanted; the others are typically discarded. Although few couples are likely to undergo the difficulty and expense of IVF simply to choose the sex of their child, embryo screening is a highly reliable means of sex selection. And as our genetic knowledge increases, it may be possible to use PGD to cull embryos carrying undesired genes, such as those associated with obesity, height, and skin color. The science-fiction movie *Gattaca* depicts a future in which parents routinely screen embryos for sex, height, immunity to disease, and even IQ. There is something troubling about the *Gattaca* scenario, but it is not easy to identify what exactly is wrong with screening embryos to choose the sex of our children.

One line of objection draws on arguments familiar from the abortion debate. Those who believe that an embryo is a person reject embryo screening for the same reasons they reject abortion. If an eight-cell embryo growing in a petri dish is morally equivalent to a fully developed human being, then discarding it is no better than aborting a fetus, and both practices are equivalent to infanticide. Whatever its merits, however, this "pro-life" objection is not an argument against sex selection as such.

The latest technology poses the question of sex selection unclouded by the matter of an embryo's moral status. The Genetics & IVF Institute, a for-profit infertility clinic in Fairfax, Virginia, now offers a sperm-sorting technique that makes it possible to choose the sex of one's child before it is conceived. X-bearing sperm, which produce girls, carry more DNA than Y-bearing sperm, which produce boys; a device called a flow cytometer can separate them. The process, called MicroSort, has a high rate of success.

If sex selection by sperm sorting is objectionable, it must be for reasons that go beyond the

debate about the moral status of the embryo. One such reason is that sex selection is an instrument of sex discrimination—typically against girls, as illustrated by the chilling sex ratios in India and China. Some speculate that societies with substantially more men than women will be less stable, more violent, and more prone to crime or war. These are legitimate worries—but the sperm-sorting company has a clever way of addressing them. It offers MicroSort only to couples who want to choose the sex of a child for purposes of "family balancing." Those with more sons than daughters may choose a girl, and vice versa. But customers may not use the technology to stock up on children of the same sex, or even to choose the sex of their firstborn child. (So far the majority of MicroSort clients have chosen girls.) Under restrictions of this kind, do any ethical issues remain that should give us pause?

The case of MicroSort helps us isolate the moral objections that would persist if muscle-enhancement, memory-enhancement, and height-enhancement technologies were safe and available to all.

It is commonly said that genetic enhancements undermine our humanity by threatening our capacity to act freely, to succeed by our own efforts, and to consider ourselves responsible-worthy of praise or blame—for the things we do and for the way we are. It is one thing to hit seventy home runs as the result of disciplined training and effort, and something else, something less, to hit them with the help of steroids or genetically enhanced muscles. Of course, the roles of effort and enhancement will be a matter of degree. But as the role of enhancement increases, our admiration for the achievement fades—or, rather, our admiration for the achievements shifts from the player to his pharmacist. This suggests that our moral response to enhancement is a response to the diminished agency of the person whose achievement is enhanced.

Though there is much to be said for this argument, I do not think the main problem with enhancement and genetic engineering is that they undermine effort and erode human agency. The deeper danger is that they represent a kind of hyperagency—a Promethean aspiration to remake nature, including human nature, to serve our purposes and satisfy our desires. The problem is not the drift to mechanism but the drive to mastery. And what the drive to mastery misses and may even destroy is an appreciation of the gifted character of human powers and achievements.

To acknowledge the giftedness of life is to recognize that our talents and powers are not wholly our own doing, despite the effort we expend to develop and to exercise them. It is also to recognize that not everything in the world is open to whatever use we may desire or devise. Appreciating the gifted quality of life constrains the Promethean project and conduces to a certain humility. It is in part a religious sensibility. But its resonance reaches beyond religion.

It is difficult to account for what we admire about human activity and achievement without drawing upon some version of this idea. Consider two types of athletic achievement. We appreciate players like Pete Rose, who are not blessed with great natural gifts but who manage, through striving, grit, and determination, to excel in their sport. But we also admire players like Joe DiMaggio, who display natural gifts with grace and effortlessness. Now, suppose we learned that both players took performance-enhancing drugs. Whose turn to drugs would we find more deeply disillusioning? Which aspect of the athletic ideal—effort or gift—would be more deeply offended?

Some might say effort: the problem with drugs is that they provide a shortcut, a way to win without striving. But striving is not the point of sports; excellence is. And excellence consists at least partly in the display of natural talents and gifts that are no doing of the athlete who possesses them. This is an uncomfortable fact for democratic societies. We want to believe that success, in sports and in life, is something we earn, not something we inherit. Natural gifts, and the admiration they inspire, embarrass the meritocratic faith; they cast doubt on the conviction that praise and rewards flow from effort alone. In the face of this embarrassment we inflate the moral significance of striving, and depreciate giftedness. This distortion can be seen, for example, in network-television coverage of the Olympics, which focuses less on the feats

the athletes perform than on heartrending stories of the hardships they have overcome and the struggles they have waged to triumph over an injury or a difficult upbringing or political turmoil in their native land.

But effort isn't everything. No one believes that a mediocre basketball player who works and trains even harder than Michael Jordan deserves greater acclaim or a bigger contract. The real problem with genetically altered athletes is that they corrupt athletic competition as a human activity that honors the cultivation and display of natural talents. From this standpoint, enhancement can be seen as the ultimate expression of the ethic of effort and willfulness—a kind of high-tech striving. The ethic of willfulness and the biotechnological powers it now enlists are arrayed against the claims of giftedness.

The ethic of giftedness, under siege in sports, persists in the practice of parenting. But here, too, bioengineering and genetic enhancement threaten to dislodge it. To appreciate children as gifts is to accept them as they come, not as objects of our design or products of our will or instruments of our ambition. Parental love is not contingent on the talents and attributes a child happens to have. We choose our friends and spouses at least partly on the basis of qualities we find attractive. But we do not choose our children. Their qualities are unpredictable, and even the most conscientious parents cannot be held wholly responsible for the kind of children they have. That is why parenthood, more than other human relationships, teaches what the theologian William F. May calls an "openness to the unbidden."

May's resonant phrase helps us see that the deepest moral objection to enhancement lies less in the perfection it seeks than in the human disposition it expresses and promotes. The problem is not that parents usurp the autonomy of a child they design. The problem lies in the hubris of the designing parents, in their drive to master the mystery of birth. Even if this disposition did not make parents tyrants to their children, it would disfigure the relation between parent and child, and deprive the parent of the humility and enlarged human sympathies that an openness to the unbidden can cultivate.

To appreciate children as gifts or blessings is not, of course, to be passive in the face of illness or disease. Medical intervention to cure or prevent illness or restore the injured to health does not desecrate nature but honors it. Healing sickness or injury does not override a child's natural capacities but permits them to flourish.

Nor does the sense of life as a gift mean that parents must shrink from shaping and directing the development of their child. Just as athletes and artists have an obligation to cultivate their talents, so parents have an obligation to cultivate their children, to help them discover and develop their talents and gifts. As May points out, parents give their children two kinds of love: accepting love and transforming love. Accepting love affirms the being of the child, whereas transforming love seeks the well-being of the child. Each aspect corrects the excesses of the other, he writes: "Attachment becomes too quietistic if it slackens into mere acceptance of the child as he is." Parents have a duty to promote their children's excellence.

These days, however, overly ambitious parents are prone to get carried away with transforming love—promoting and demanding all manner of accomplishments from their children, seeking perfection. "Parents find it difficult to maintain an equilibrium between the two sides of love," May observes. "Accepting love, without transforming love, slides into indulgence and finally neglect. Transforming love, without accepting love, badgers and finally rejects." May finds in these competing impulses a parallel with modern science: it, too, engages us in beholding the given world, studying and savoring it, and also in molding the world, transforming and perfecting it.

The mandate to mold our children, to cultivate and improve them, complicates the case against enhancement. We usually admire parents who seek the best for their children, who spare no effort to help them achieve happiness and success. Some parents confer advantages on their children by enrolling them in expensive schools, hiring private tutors, sending them to tennis camp, providing them with piano lessons, ballet lessons, swimming lessons, SAT-prep courses, and so on. If it is permissible and even admirable for parents to help their children in these ways, why isn't

it equally admirable for parents to use whatever genetic technologies may emerge (provided they are safe) to enhance their children's intelligence, musical ability, or athletic prowess?

The defenders of enhancement are right to this extent: improving children through genetic engineering is similar in spirit to the heavily managed, high-pressure child-rearing that is now common. But this similarity does not vindicate genetic enhancement. On the contrary, it highlights a problem with the trend toward hyperparenting. One conspicuous example of this trend is sports-crazed parents bent on making champions of their children. Another is the frenzied drive of overbearing parents to mold and manage their children's academic careers.

As the pressure for performance increases, so does the need to help distractible children concentrate on the task at hand. This may be why diagnoses of attention deficit and hyperactivity disorder have increased so sharply. Lawrence Diller, a pediatrician and the author of *Running on Ritalin*, estimates that five to six percent of American children under eighteen (a total of four to five million kids) are currently prescribed Ritalin, Adderall, and other stimulants, the treatment of choice for ADHD. (Stimulants counteract hyperactivity by making it easier to focus and sustain attention.) The number of Ritalin prescriptions for children and adolescents has tripled over the past decade, but not all users suffer from attention disorders or hyperactivity. High school and college students have learned that prescription stimulants improve concentration for those with normal attention spans, and some buy or borrow their classmates' drugs to enhance their performance on the SAT or other exams. Since stimulants work for both medical and nonmedical purposes, they raise the same moral questions posed by other technologies of enhancement.

However those questions are resolved, the debate reveals the cultural distance we have traveled since the debate over marijuana, LSD, and other drugs a generation ago. Unlike the drugs of the 1960s and 1970s Ritalin and Adderall are not for checking out but for buckling down, not for beholding the world and taking it in but for molding the world and fitting in. We used to speak of nonmedical drug use as "recreational." That term no longer applies. The steroids and stimulants that figure in the enhancement debate are not a source of recreation but a bid for compliance—a way of answering a competitive society's demand to improve our performance and perfect our nature. This demand for performance and perfection animates the impulse to rail against the given. It is the deepest source of the moral trouble with enhancement.

Some see a clear line between genetic enhancement and other ways that people seek improvement in their children and themselves. Genetic manipulation seems somehow worse—more intrusive, more sinister—than other ways of enhancing performance and seeking success. But morally speaking, the difference is less significant than it seems. Bioengineering gives us reason to question the low-tech, high-pressure child-rearing practices we commonly accept. The hyper-parenting familiar in our time represents an anxious excess of mastery and dominion that misses the sense of life as a gift. This draws it disturbingly close to eugenics.

The shadow of eugenics hangs over today's debates about genetic engineering and enhancement. Critics of genetic engineering argue that human cloning, enhancement, and the quest for designer children are nothing more than "privatized" or "free-market" eugenics. Defenders of enhancement reply that genetic choices freely made are not really eugenic—at least not in the pejorative sense. To remove the coercion, they argue, is to remove the very thing that makes eugenic policies repugnant.

Sorting out the lesson of eugenics is another way of wrestling with the ethics of enhancement. The Nazis gave eugenics a bad name. But what, precisely, was wrong with it? Was the old eugenics objectionable only insofar as it was coercive? Or is there something inherently wrong with the resolve to deliberately design our progeny's traits?

James Watson, the biologist who, with Francis Crick, discovered the structure of DNA, sees nothing wrong with genetic engineering and enhancement, provided they are freely chosen rather than state-imposed. And yet Watson's

language contains more than a whiff of the old eugenic sensibility. "If you really are stupid, I would call that a disease," he recently told *The Times of London*. "The lower 10 percent who really have difficulty, even in elementary school, what's the cause of it? A lot of people would like to say, 'Well, poverty, things like that.' It probably isn't. So I'd like to get rid of that, to help the lower 10 percent." A few years ago Watson stirred controversy by saying that if a gene for homosexuality were discovered, a woman should be free to abort a fetus that carried it. When his remark provoked an uproar, he replied that he was not singling out gays but asserting a principle: women should be free to abort fetuses for any reason of genetic preference—for example, if the child would be dyslexic, or lacking musical talent, or too short to play basketball.

Watson's scenarios are clearly objectionable to those for whom all abortion is an unspeakable crime. But for those who do not subscribe to the pro-life position, these scenarios raise a hard question: If it is morally troubling to contemplate abortion to avoid a gay child or a dyslexic one, doesn't this suggest that something is wrong with acting on any eugenic preference, even when no state coercion is involved?

Consider the market in eggs and sperm. The advent of artificial insemination allows prospective parents to shop for gametes with the genetic traits they desire in their off-spring. It is a less predictable way to design children than cloning or preimplantation genetic screening, but it offers a good example of a procreative practice in which the old eugenics meets the new consumerism. A few years ago some Ivy League newspapers ran an ad seeking an egg from a woman who was at least five feet ten inches tall and athletic, had no major family medical problems, and had a combined SAT score of 1400 or above. The ad offered $50,000 for an egg from a donor with these traits. More recently a Web site was launched claiming to auction eggs from fashion models whose photos appeared on the site, at starting bids of $15,000 to $150,000.

On what grounds, if any, is the egg market morally objectionable? Since no one is forced to buy or sell, it cannot be wrong for reasons of coercion. Some might worry that hefty prices would exploit poor women by presenting them with an offer they couldn't refuse. But the designer eggs that fetch the highest prices are likely to be sought from the privileged, not the poor. If the market for premium eggs gives us moral qualms, this, too, shows that concerns about eugenics are not put to rest by freedom of choice.

A tale of two sperm banks helps explain why. The Repository for Germinal Choice, one of America's first sperm banks, was not a commercial enterprise. It was opened in 1980 by Robert Graham, a philanthropist dedicated to improving the world's "germ plasm" and counteracting the rise of "retrograde humans." His plan was to collect the sperm of Nobel Prize–winning scientists and make it available to women of high intelligence, in hopes of breeding super-smart babies. But Graham had trouble persuading Nobel laureates to donate their sperm for his bizarre scheme, and so settled for sperm from young scientists of high promise. His sperm bank closed in 1999.

In contrast, California Cryobank, one of the world's leading sperm banks, is a for-profit company with no overt eugenic mission. Cappy Rothman, M.D., a co-founder of the firm, has nothing but disdain for Graham's eugenics, although the standards Cryobank imposes on the sperm it recruits are exacting. Cryobank has offices in Cambridge, Massachusetts, between Harvard and MIT, and in Palo Alto, California, near Stanford. It advertises for donors in campus newspapers (compensation up to $900 a month), and accepts less than five percent of the men who apply. Cryobank's marketing materials play up the prestigious source of its sperm. Its catalog provides detailed information about the physical characteristics of each donor, along with his ethnic origin and college major. For an extra fee prospective customers can buy the results of a test that assesses the donor's temperament and character type. Rothman reports that Cryobank's ideal sperm donor is six feet tall, with brown eyes, blond hair, and dimples, and has a college degree—not because the company wants to propagate those traits, but because those are the traits

his customers want: "If our customers wanted high school dropouts, we would give them high school dropouts."

Not everyone objects to marketing sperm. But anyone who is troubled by the eugenic aspect of the Nobel Prize sperm bank should be equally troubled by Cryobank, consumer-driven though it be. What, after all, is the moral difference between designing children according to an explicit eugenic purpose and designing children according to the dictates of the market? Whether the aim is to improve humanity's "germ plasm" or to cater to consumer preferences, both practices are eugenic insofar as both make children into products of deliberate design.

A number of political philosophers call for a new "liberal eugenics." They argue that a moral distinction can be drawn between the old eugenic policies and genetic enhancements that do not restrict the autonomy of the child. "While old-fashioned authoritarian eugenicists sought to produce citizens out of a single centrally designed mould," writes Nicholas Agar, "the distinguishing mark of the new liberal eugenics is state neutrality." Government may not tell parents what sort of children to design, and parents may engineer in their children only those traits that improve their capacities without biasing their choice of life plans. A recent text on genetics and justice, written by the bioethicists Allen Buchanan, Dan W. Brock, Norman Daniels, and Daniel Wilder, offers a similar view. The "bad reputation of eugenics," they write, is due to practices that "might be avoidable in a future eugenic program." The problem with the old eugenics was that its burdens fell disproportionately on the weak and the poor, who were unjustly sterilized and segregated. But provided that the benefits and burdens of genetic improvement are fairly distributed, these bioethicists argue, eugenic measures are unobjectionable and may even be morally required.

The libertarian philosopher Robert Nozick proposed a "genetic supermarket" that would enable parents to order children by design without imposing a single design on the society as a whole: "This supermarket system has the great virtue that it involves no centralized decision fixing the future human type(s)."

Even the leading philosopher of American liberalism, John Rawls, in his classic *A Theory of Justice* (1971), offered a brief endorsement of noncoercive eugenics. Even in a society that agrees to share the benefits and burdens of the genetic lottery, it is "in the interest of each to have greater natural assets," Rawls wrote. "This enables him to pursue a preferred plan of life." The parties to the social contract "want to insure for their descendants the best genetic endowment (assuming their own to be fixed)." Eugenic policies are therefore not only permissible but required as a matter of justice. "Thus over time a society is to take steps at least to preserve the general level of natural abilities and to prevent the diffusion of serious defects."

But removing the coercion does not vindicate eugenics. The problem with eugenics and genetic engineering is that they represent the one-sided triumph of willfulness over giftedness, of dominion over reverence, of molding over beholding. Why, we may wonder, should we worry about this triumph? Why not shake off our unease about genetic enhancement as so much superstition? What would be lost if biotechnology dissolved our sense of giftedness?

From a religious standpoint the answer is clear. To believe that our talents and powers are wholly our own doing is to misunderstand our place in creation, to confuse our role with God's. Religion is not the only source of reasons to care about giftedness, however. The moral stakes can also be described in secular terms. If bioengineering made the myth of the "self-made man" come true, it would be difficult to view our talents as gifts for which we are indebted, rather than as achievements for which we are responsible. This would transform three key features of our moral landscape: humility, responsibility, and solidarity.

In a social world that prizes mastery and control, parenthood is a school for humility. That we care deeply about our children and yet cannot choose the kind we want teaches parents to be open to the unbidden. Such openness is a disposition worth affirming, not only within

families but in the wider world as well. It invites us to abide the unexpected, to live with dissonance, to rein in the impulse to control. A *Gattaca*-like world in which parents became accustomed to specifying the sex and genetic traits of their children would be a world inhospitable to the unbidden, a gated community writ large. The awareness that our talents and abilities are not wholly our own doing restrains our tendency toward hubris.

Though some maintain that genetic enhancement erodes human agency by overriding effort, the real problem is the explosion, not the erosion, of responsibility. As humility gives way, responsibility expands to daunting proportions. We attribute less to chance and more to choice. Parents become responsible for choosing, or failing to choose, the right traits for their children. Athletes become responsible for acquiring, or failing to acquire, the talents that will help their teams win.

One of the blessings of seeing ourselves as creatures of nature, God, or fortune is that we are not wholly responsible for the way we are. The more we become masters of our genetic endowments, the greater the burden we bear for the talents we have and the way we perform. Today when a basketball player misses a rebound, his coach can blame him for being out of position. Tomorrow the coach may blame him for being too short. Even now the use of performance-enhancing drugs in professional sports is subtly transforming the expectations players have for one another; on some teams players who take the field free from amphetamines or other stimulants are criticized for "playing naked."

The more alive we are to the chanced nature of our lot, the more reason we have to share our fate with others. Consider insurance. Since people do not know whether or when various ills will befall them, they pool their risk by buying health insurance and life insurance. As life plays itself out, the healthy wind up subsidizing the unhealthy, and those who live to a ripe old age wind up subsidizing the families of those who die before their time. Even without a sense of mutual obligation, people pool their risks and resources and share one another's fate.

But insurance markets mimic solidarity only insofar as people do not know or control their own risk factors. Suppose genetic testing advanced to the point where it could reliably predict each person's medical future and life expectancy. Those confident of good health and long life would opt out of the pool, causing other people's premiums to skyrocket. The solidarity of insurance would disappear as those with good genes fled the actuarial company of those with bad ones.

The fear that insurance companies would use genetic data to assess risks and set premiums recently led the Senate to vote to prohibit genetic discrimination in health insurance. But the bigger danger, admittedly more speculative, is that genetic enhancement, if routinely practiced, would make it harder to foster the moral sentiments that social solidarity requires.

Why, after all, do the successful owe anything to the least-advantaged members of society? The best answer to this question leans heavily on the notion of giftedness. The natural talents that enable the successful to flourish are not their own doing but, rather, their good fortune—a result of the genetic lottery. If our genetic endowments are gifts, rather than achievements for which we can claim credit, it is a mistake and a conceit to assume that we are entitled to the full measure of the bounty they reap in a market economy. We therefore have an obligation to share this bounty with those who, through no fault of their own, lack comparable gifts.

A lively sense of the contingency of our gifts—a consciousness that none of us is wholly responsible for his or her success—saves a meritocratic society from sliding into the smug assumption that the rich are rich because they are more deserving than the poor. Without this, the successful would become even more likely than they are now to view themselves as self-made and self-sufficient, and hence wholly responsible for their success. Those at the bottom of society would be viewed not as disadvantaged, and thus worthy of a measure of

compensation, but as simply unfit, and thus worthy of eugenic repair. The meritocracy, less chastened by chance, would become harder, less forgiving. As perfect genetic knowledge would end the simulacrum of solidarity in insurance markets, so perfect genetic control would erode the actual solidarity that arises when men and women reflect on the contingency of their talents and fortunes.

Thirty-five years ago Robert L. Sinsheimer, a molecular biologist at the California Institute of Technology, glimpsed the shape of things to come. In an article titled "The Prospect of Designed Genetic Change" he argued that freedom of choice would vindicate the new genetics, and set it apart from the discredited eugenics of old.

> To implement the older eugenics would have required a massive social program carried out over many generations. Such a program could not have been initiated without the consent and cooperation of a major fraction of the population, and would have been continuously subject to social control. In contrast, the new eugenics could, at least in principle, be implemented on a quite individual basis, in one generation, and subject to no existing restrictions.

According to Sinsheimer, the new eugenics would be voluntary rather than coerced, and also more humane. Rather than segregating and eliminating the unfit, it would improve them. "The old eugenics would have required a continual selection for breeding of the fit, and a culling of the unfit," he wrote. "The new eugenics would permit in principle the conversion of all the unfit to the highest genetic level."

Sinsheimer's paean to genetic engineering caught the heady, Promethean self-image of the age. He wrote hopefully of rescuing "the losers in that chromosomal lottery that so firmly channels our human destinies," including not only those born with genetic defects but also "the 50,000,000 'normal' Americans with an IQ of less than 90." But he also saw that something bigger than improving on nature's "mindless, age-old throw of dice" was at stake. Implicit in technologies of genetic intervention was a more exalted place for human beings in the cosmos. "As we enlarge man's freedom, we diminish his constraints and that which he must accept as given," he wrote, Copernicus and Darwin had "demoted man from his bright glory at the focal point of the universe," but the new biology would restore his central role. In the mirror of our genetic knowledge we would see ourselves as more than a link in the chain of evolution: "We can be the agent of transition to a whole new pitch of evolution. This is a cosmic event."

There is something appealing, even intoxicating, about a vision of human freedom unfettered by the given. It may even be the case that the allure of that vision played a part in summoning the genomic age into being. It is often assumed that the powers of enhancement we now possess arose as an inadvertent by-product of biomedical progress—the genetic revolution came, so to speak, to cure disease, and stayed to tempt us with the prospect of enhancing our performance, designing our children, and perfecting our nature. That may have the story backwards. It is more plausible to view genetic engineering as the ultimate expression of our resolve to see ourselves astride the world, the masters of our nature. But that promise of mastery is flawed. It threatens to banish our appreciation of life as a gift, and to leave us with nothing to affirm or behold outside our own will.

2.6 ENERGY AND THE ENVIRONMENT

2.6.1 THE NEXUS AND THE NEEM TREE

ROBERT W. KATES

In this reading, Robert W. Kates, university professor emeritus at Brown University and coconvener of the Initiative on Science and Technology for Sustainability, examines the concept of sustainable development in light of what he regards as the failures of the contemporary pattern of globalization. Kates characterizes sustainable development as an oxymoron that "seeks to finesse the real conflicts between economy and environment and between the present and the future." Finding the correct balance between these ends requires that we examine the current system of globalization to determine whether or not it is enabling us to meet the goals of reducing the numbers of chronically hungry people while preserving Earth's life-support system for future generations. On both of these counts, Kates believes that the current form of globalization is failing to meet these goals overall, even though there are important and hopeful signs of progress in some areas. Instead of being opposed to globalization, the challenge, as he sees it, is to "humanize" it by changing the way in which the current system of globalization is governed in the direction of greater local control and democracy at the grass-roots level. Only by "civilizing" globalization in this way can we ensure that the benefits of economic development will be fairly distributed and used to address human needs.

FOCUS QUESTIONS

1. How does the U.S. National Academy of Sciences—National Research Council define the concept of sustainable development? To what extent are the basic human needs identified in this concept adequately satisfied for the approximately 6 billion humans living today? Discuss.
2. In what respects does the Neem tree stand as a metaphor for globalization? What are the defining characteristics of globalization? Has globalization, as it has been developed thus far, helped to reduce the number of chronically hungry people in the world? Why or why not?
3. What does the IPAT formula tell us about the causes of environmental degradation? Why does the author believe that this formula is only partially accurate? What factors need to be added to improve it? Explain.

KEYWORDS

consumption, globalization, governance, population, risk assessment, sustainable development

Source: Robert Kates, "The Nexus and the Neem Tree," from *Worlds Apart* edited by James Gustave Speth, pp. 85–107. Copyright © 2003 Island Press. Reproduced by permission of Island Press, Washington, DC.

The world of this new century is in transition—becoming more crowded and more consuming, warmer and more stressed, more interconnected, yet diverse and divided. Can this transition also be a transition for sustainability, in which the more than 9 billion people of the next half century meet their wants and needs in ways that do not further degrade the planet's life-support systems? In this chapter, I explore this transition to sustainability, the context in which it will take place, and the ways in which the new-old phenomenon of globalization affects it. My title is a play on the title of what may be the best-known book on globalization—Tom Friedman's *The Lexus and the Olive Tree*.

Friedman describes his title as follows:

> So there I was speeding along at 180 miles an hour on the most modern train in the world, reading this story about the oldest corner of the world. And the thought occurred to me that these Japanese, whose Lexus factory I had just visited and whose train I was riding, were building the greatest luxury car in the world with robots. And over here, on the top of page 3 of the *Herald Tribune*, the people with whom I had lived for so many years in Beirut and Jerusalem, whom I knew so well, were still fighting over who owned which olive tree. It struck me then that the Lexus and the olive tree were pretty good symbols of this post-Cold War era: half the world seemed to be emerging from the Cold War intent on building a better Lexus, dedicated to modernizing, streamlining and privatizing their economies in order to thrive in the system of globalization. And half the world—sometimes half the same country, sometimes half the same person—was still caught up in the fight over who owns which olive tree (p. 31).

But unlike this Lexus, which stands for all that is modern and different in globalization, and the olive tree for all that resists it, the nexus I explore is that of environment and development and the context in which it will play out over the coming two generations. And the Neem tree symbolizes the globalizing world where the elements of both the Lexus and the olive tree coexist in an uneasy tension of mutual attraction and repulsion.

I use three critical goals required for a successful sustainability transition—meeting human needs, reducing hunger and poverty, and preserving the life-support systems of the planet—to ask how globalization might help or hinder achieving these. For globalization to help more than hinder, it will need to be "civilized," and I conclude with an analog of how that might take place.

SUSTAINABLE DEVELOPMENT: THE NEXUS OF ENVIRONMENT AND DEVELOPMENT

The nexus of society's developmental goals with its environmental limits over the long term comes together in "sustainable development," which is only the most recent effort to link together the collective aspirations of the peoples of the world. Over my adult life, four aspirations emerged: first, for peace in the postwar world of 1945; then for freedom, in the struggles in the late 1940s and 1950s to end imperialism; followed by development for the poorest three-fourths of the world; and last, in the final quarter of the century, a concern for a healthy environment for humankind, the earth itself, and its complex systems that support life. As global aspirations develop, good people try to bring them together in a characteristic pattern of international high-level commissions (Brandt, Palme, Brundtland), followed by great international conferences. Such was the 1987 report of the World Commission on Environment and Development (WCED, also known as the Brundtland Commission) widely disseminated as *Our Common Future* (WCED, 1987), followed by the United Nations Conference on Environment and Development (UNCED) in Rio de Janeiro in 1992, and now a decade later in South Africa as the World Summit for Sustainable Development.

"Sustainable development" is now central to the mission of countless international organizations, national institutions, "sustainable cities" and locales, transnational corporations, and nongovernmental organizations. The genius of the oxymoron of sustainable development lies in its essential ambiguity that seeks to finesse the real conflicts between economy and environment and between the present and the future. While sharing

a common concern for the fate of the earth, proponents of sustainable development differ in their emphases on what is to be sustained, what is to be developed, how to link environment and development, and for how long a time (see Figure 1). Thus, proponents differ on what is to be sustained: Is it nature itself, or nature in the utilitarian life support of humankind, and does it include sustaining the community of the olive tree, as Friedman (2000) describes it: "everything that roots us ... family, a community, a tribe, a nation, a religion or, most of all, a place called home" (p. 31)? Proponents differ on what is to be developed: Is it the economy, some broader notion of society, or is it people themselves? And how shall we link the two: sustain only, develop mostly, develop only but sustain somewhat, sustain or develop?—these and many more permutations are found.

Finally, over what time horizon will this occur? The Brundtland report employs the usefully ambiguous and widely accepted time horizon as "now and in the future." But in a future of a single generation, twenty-five years, almost any development appears sustainable. Over an infinite forever, none does, as even the smallest growth extended indefinitely creates situations that seem surely unsustainable. And over the century, now encompassed in many assessments such as that of climate change, the large and the long future is both remote and uncertain.

While a major political success, sustainable development has not been a significant scientific focus beyond the earliest days of its conceptualization. While originating in the scientific activities of the early 1980s, particularly the work of the International Union for the Conservation of Nature (IUCN), as sustainable development gained greater political adherence, organized science found less to address. This has now changed, and a focus on a transition toward sustainability has made sustainable development scientifically manageable, and measurable for the world academies of science (IAP, 2000), for the international organizations of science (ICSU/WFEO, 2002), and increasingly for an emergent sustainability science (Kates et al., 2001).

In 1995, the Board on Sustainable Development of the U.S. National Academy of Sciences-National Research Council (NAS-NRC) sought to resolve some of the ambiguity I have described by focusing on a transition to sustainability over the next fifty years. By focusing on a transition,

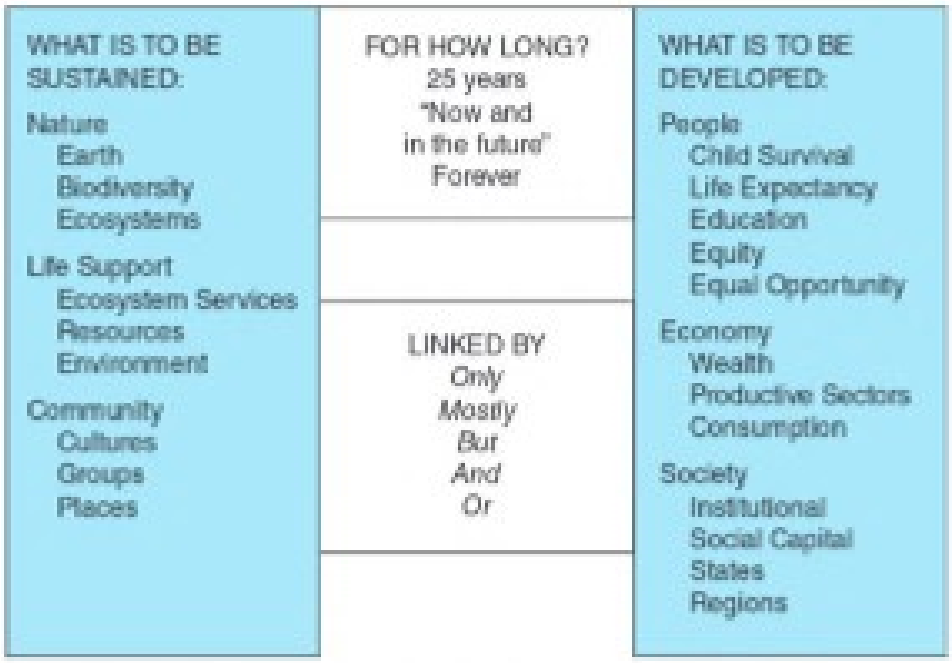

FIGURE 1 Sustainable Development: Common Concerns, Differing Emphases.

we took as our starting point the best understood of future trends—the demographic transition from a world with a population that grew by many births and many deaths to one that stabilized with few births and few deaths. With such a transition well under way, we can with some confidence project a declining or steady state population by the end of this century, with the bulk of that population born by 2050 (perhaps 9 billion of an eventual 10 billion). Thus, the human development needs of that population will surely increase, compared to the more than 6 billion alive today, but probably not more than by half again as much as that of today. The board defined a sustainability transition as one that would meet the human needs for food, nurture, housing, education, and employment of that larger but finite population, significantly reducing hunger and poverty, while still maintaining the essential life-support systems of the planet (NRC-BSD, 1999). These human needs are unmet today; for example, in 1995, 16 percent of the world population was hungry, 24 percent had unsafe water to drink, and 24 percent were illiterate (Raskin et al., 1998).

For these three normative goals, we found ample consensual support and measurable targets in the deliberations and subsequent treaties of international conferences and summits of leaders. For example, for the amount of reduction in hunger and poverty we used international consensus statements that call for reducing hunger and poverty by half within one to two decades (IMF et al., 2000) and suggested a target of reducing hunger by half in each of the next two generations.

Compared to meeting human needs, quantitative targets for preserving life-support systems are fewer, more modest, and more contested. Global targets now exist for reducing ozone-depleting substances, greenhouse gases, and, regionally, for some air pollutants. Absolute prohibitions (zero targets) exist for ocean dumping of radioactive wastes and some toxics (persistent organic pesticides), for the taking and/or sale of a few large mammals (whales, elephants, seals), migratory birds when breeding or endangered, and certain regional fishing stocks. International standards exist for many toxic materials, organic pollutants, and heavy metals that threaten human health, but not for ecosystem health. Water, land, and vegetative resources, such as arid lands or forests, have at best qualitative aspirations for sustainable management or restoration.

THE NEEM TREE: THE NEW AND OLD GLOBALIZATION

The 1992 NAS-NRC report *Neem: A Tree for Solving Global Problems* begins:

> Neem is a fascinating tree. On the one hand, it seems to be one of the most promising of all plants, and may eventually benefit every person on the planet. Probably no other yields as many strange and varied products or has as many exploitable by-products. Indeed, as foreseen by some scientists, this plant may usher in [a] new era in pest control, provide millions with inexpensive medicines, cut down the rate of human population growth, and perhaps even reduce erosion, deforestation, and the excessive temperature of an overheated globe.

The Neem tree, *Azadirachta indica,* is an attractive broad-leafed evergreen that grows tall and broad and can live for a century or more. Native to South Asia, it has been carried over the last century to the rest of tropical and semitropical Asia, Africa, and increasingly to the Caribbean and Central America and is now well established in thirty countries and has been introduced to many more. Everywhere it grows, it is prized for its ability to grow in marginal soils, to provide shade, firewood, oil for lamps, cosmetics, soaps, lubrication, and medicinals that date back several millennia, twig toothbrushes that prevent gum disease, and as a natural insecticide. Even where it does not grow, word of its wonders [is] carried on numerous Web sites, many dedicated to the Neem itself.

But it is for its pesticidal qualities—as a safer alternative to dangerous neurotoxins, effective across a large range of insects, fungi, nematodes, and the like, and seemingly safe for humans, birds, and animals—that Neem has attracted considerable scientific and commercial interest. More than seventy patents for uses or processes related

to Neem products exist, and in May 2000 in an important decision, the European Patent Office revoked the patent given to W. R. Grace company for a fungicidal product, a decision hailed by Vandana Shiva, who had challenged the original patent, as "a great day for all who have been fighting to take back control of their resources and knowledge-systems from the patent regimes of the North" (Anon., 2000).

Thus, the Neem shares three major characteristics of globalization. It is not new, but quite ancient, as is globalization. In ancient Sanskrit it is known as *aristha*, or reliever of sickness. As with previous globalizations, it spread with religion and with empire. Part of traditional ayurvedic medicine, the Neem is found wherever Hinduism is found and often where the British Empire ruled. But as with globalization, there is much that is new in its dispersal and product development, as Neem seeds are now an international commodity. Modern science is close to synthesizing its major insecticidal properties, international nongovernmental organizations encourage its usage for impoverished rural peoples, and Web sites huckster its cosmetic and medicinal values. Finally, as with globalization, it is full of unrealized promise and currently realized discord. For except as a source of shade, firewood, toothbrushes, oil, or home remedies, it is not widely used beyond its South Asian home; and its most promising commercial products—refined or synthesized and standardized pesticides, medicinals, or contraceptives—are either underdeveloped or contested as to their efficacy and safety, as well as ethically in relation to the commercialization of an ancient legacy of nature and humankind.

CONTEMPORARY GLOBALIZATION

As to globalization in general, I prefer the simplified definition by Held, McGrew, Goldblatt, and Perraton (1999) in *Global Transformations: Politics, Economics, and Culture*, a study that some think is currently the best academic book on the subject. Held et al. say that "in its simplest sense globalization refers to the widening, deepening and speeding up of global interconnectedness..." (p. 14). But, of course, as good academics they are not content with such simplicity and go on to describe a set of technical terms and criteria to mirror these items as extensiveness, intensiveness, velocity, and the impacts of interconnections.

Globalization, as noted, is not new, and Held et al. recognize four major periods of globalization: the premodern period of early empires and world religions, the early modern period of Western expansion, the modern industrial era, and the contemporary period from 1945 to the present. I would add two others: the earliest prehistoric period in which humans spread out of Africa around the world, and the future, especially that of the first half of the twenty-first century.

Reviewing the contemporary period and projecting to the future, our academy study (NRC-BSD, 1999) identified some major dimensions of contemporary globalization. The first is global interconnectedness with the much larger population of the future more closely connected by ties of economic production and consumption, migration, communication, and interlinked technologies. Since 1950, trade between nations has grown at more than twice the rate of the economy, and now some 20 percent of the world's goods and services pass over a border. Trade in money and capital—a hundred times the volume of world trade—now moves at a dizzying pace with electronic movement of funds, worldwide currency markets, and twenty-four-hour financial markets.

Words, images, and ideas also outpace the flow of products. New information technologies and mass communication techniques will continue to penetrate many different linguistic, cultural, and political barriers. Flows of people—temporary, permanent, and forced—have also increased, although most movements are poorly measured. The rate of increase in refugees is more rapid than that of world trade.

The rapid movement of peoples and products also makes possible the rapid transmission of infectious diseases of people, crops, and livestock and the biological invasions so destructive of native biota. Environmental harms are exported to countries with weak environmental standards. Most feared of all may be the rapid increases in consumption fueled by aggressive marketing and

rapid cultural change. But as communication carries a culture of consumption, it also carries a culture of universal concern with the fate of the earth and links to common international efforts, shared information, and growing numbers of environmental groups.

But the academy study also considered the persistence of diversity, how connectedness, while increasing the similarity of places, can also increase diversity, particularly in urban areas that attract migrants. Places of wealth or opportunity toward which people and products are drawn actually become more diverse. There are also strong countercurrents to global culture that emphasize ethnic, national, and religious distinctiveness.

Finally, connectedness and diversity are also reflected in institutional innovation and power shifts (Mathews, 1997). At a global level, new institutions of governance have emerged, transnational corporate and financial institutions grow and consolidate, and networks of nongovernmental institutions collaborate and expand. At the subnational level, government has devolved, privatization is common, and civic society in many places has been strengthened. Power has shifted from the national state—upward to the global level and downward to the local level—and at all levels from the public to the private.

How does globalization affect a transition toward sustainability in meeting human needs, reducing hunger and poverty, and preserving life-support systems? It helps in some ways, hinders in others, and for many important characteristics, it does both. In this exploration of a highly complex subject, I present two illustrative examples using qualitative, but ordered, judgments. For one, I combine the first two human needs, feeding and nurturing, with the related goal of reducing hunger. In the second, I explore the determinants of threats to the life-support systems.

FEEDING, NURTURING, AND REDUCING HUNGER

There are three major types of hunger: chronic household hunger, episodic hunger, and special-needs hunger. Each responds somewhat

Globalization and Reducing Chronic Hunger

Current estimates find some 800–900 million people who are chronically hungry—living in households with insufficient income or its equivalent to provide for health, children's growth, and ability to work. The numbers of hungry people differ greatly between regions of the world, with the largest numbers in Asia and the greatest proportion in Africa (Table 1).

The number of chronically hungry people can be approximately estimated by using four variables: the size of the population, the average income per person, the distribution of income across the population, and the definition of a hunger line of income, or its equivalent, below which the population is thought to be hungry (see Figure 2).

It is useful to examine some differences in these major determinants over the period of contemporary globalization beginning in 1950. World *population* grew over a half century from 2.5 billion to 6 billion, but the peak growth rate was in the 1960s and has been slowing ever since. Nonetheless, about 80 million people are added each year, increasing the numbers of chronically hungry even as the proportion of hungry people

TABLE 1 Numbers and Proportion of Chronically Hungry Population, 1996–1998

REGIONS	TOTAL POPULATION (%)	HUNGRY POPULATION (MILLIONS)
Sub-Saharan Africa	34	186
Near East/ North Africa	10	36
Latin America and the Caribbean	11	55
China and India	16	348
Other Asia	19	166
Developing countries: Total	18	791

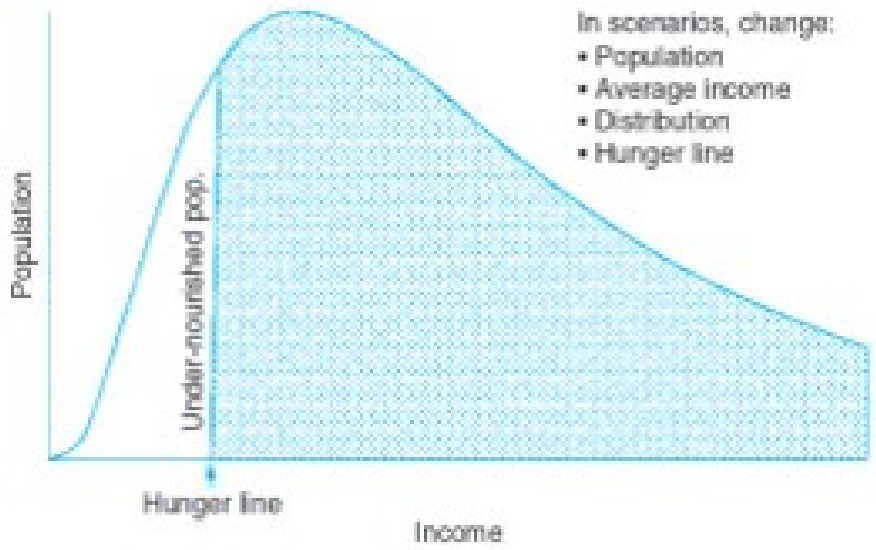

FIGURE 2 Hunger and Income.

diminishes. Examining trends in *income* (Table 2) in this period, per capita gross domestic product (GDP) grew in all world regions between 1950 and 1973, but between 1973 and 1992, GDP in Eastern Europe declined, and it stagnated in Africa and Latin America.

Inequality in the *distribution of income* occurs both between countries and regions and within countries and regions. Between regions, Africa, Eastern Europe, and Latin America (Table 2) show growing inequality with the United States (ratio of U.S. GDP per capita to regional GDP per capita), even as the rest of Europe converged with the United States and the ratio of U.S. per capita GDP to Asia declines from sevenfold in 1973 to fourfold in the course of two decades. Using a different data set (Deininger and Squire, 1996), within-region inequality of income (as measured by the Gini coefficient, a standard measure of inequality) differs almost twofold with the greatest inequality found in Latin America and the least found among the former socialist countries of Eastern Europe. Over time, for most regions, within-region inequality has been generally diminishing except in Africa and in Eastern Europe with the end of socialism. Finally, the *hunger line* grows over time as income increases, access to informal sources of food declines, food purchases increase, and diets change.

TABLE 2 Trends in Regional GDP Per Capita, 1950, 1973, and 1992

WORLD REGIONS	1950 GDP/CAPITA (IN 1990 $)	GDP/CAPITA/1950 GDP/CAPITA		U.S. GDP/CAPITA/REGIONAL GDP/CAPITA		
		1973	1992	1950	1973	1992
United States	9573	1.7	2.3	1.0	1.0	1.0
Western Europe	5513	2.1	3.2	1.7	1.4	1.2
East Europe	2235	2.4	2.0	4.3	3.1	4.7
Latin America	3478	1.4	1.7	2.8	3.3	3.6
Asia	863	2.8	6.1	11.1	6.8	4.1
Africa	893	1.5	1.5	10.7	12.5	16.2

Source: Madison, 1995.

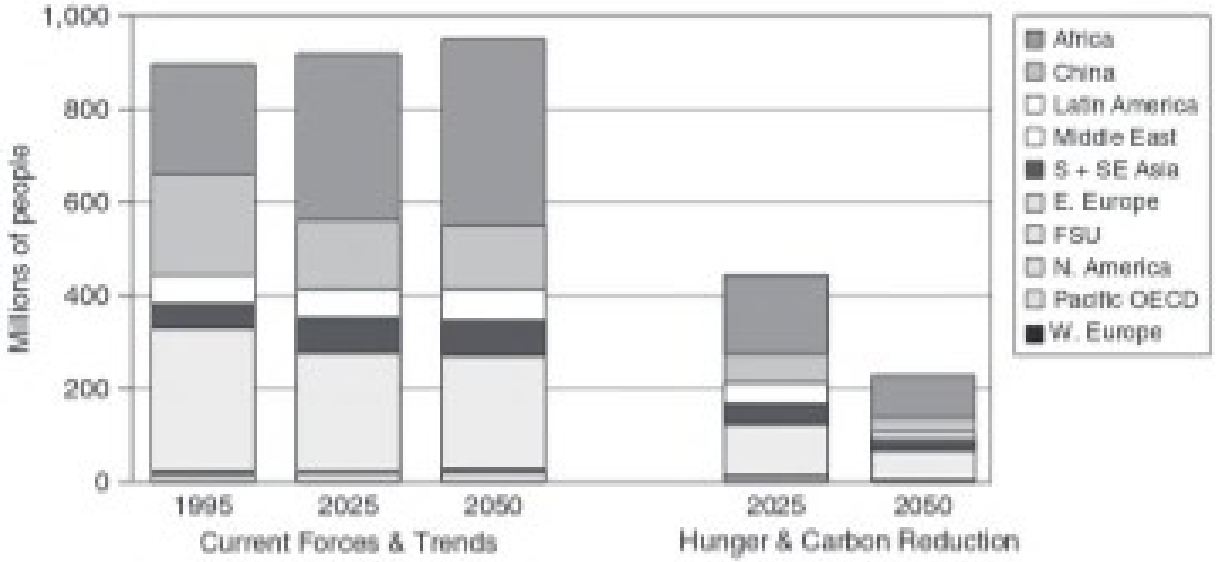

FIGURE 3 Reducing Chronic Hunger by Halves.

The NAS-NRC Board on Sustainable Development commissioned a study on the feasibility of reducing hunger by half in each of the two generations before 2050 using these variables and contrasting two different scenarios. The reference scenario projected major current trends, institutional continuity, economic globalization, and the slow convergence of developing countries toward the socioeconomic norms of developed countries. In this scenario, the number of hungry actually increased by 2050, although the proportion of hungry people declined. The "policy reform" scenario assumed that a proactive set of initiatives is instituted to achieve sustainability goals. In this scenario, hunger is cut in half with a small change in the speed of the demographic transition leading to less population, with growth in income at the higher end of plausible income growth rates, and most important, a convergence of equity to the current levels of Europe. The contrasting scenarios are shown in Figure 3 for the world as a whole and by region.

For these contrasting scenarios and the four major causal elements that underlie them, what are the impacts of globalization? In Table 3, I set out my qualitative judgments (using a scale of one plus sign [+] for small impact and four plus signs [++++] for very large impacts) as to how globalization helps or hinders the causal elements linked to chronic hunger. For *population*, globalization probably helps through a population growth rate decline by influencing all three important determinants for reducing fertility: making contraception more accessible; providing opportunities for education and work for women; and encouraging postponement of marriage through such opportunities for education and work, as well as through diffused Western lifestyle concepts.

For *income*, globalization helps increase per capita income in some parts of the world, but practices not-so-benign neglect in others. The development of an export-oriented industry in Southeast Asia, accompanied by significant public sector actions, led to major reductions in hunger in that region until the recent financial crisis

TABLE 3 Globalization's Impacts on Reducing Chronic Hunger

CAUSAL ELEMENT	HELP	HINDER
Population Growth Decrease	++	
Increase in Income per capita	++	++ In Africa
Decrease In Inequity	+ Between Nations	+++ Within Nations
Decrease in Hunger Line		+

slowed and even reversed some of those gains. But in Africa, where hunger will increase most, globalization has exacerbated some of the region's problems, its export trade in such products of affluence as oil or diamonds fueling corruption and conflict, while development aid has diminished without an equivalent growth in private investment. But most of all, it has suffered not-so-benign neglect and has been marginalized from the globalized world system.

While globalization will probably decrease *inequity* overall between countries, globalization will, for some time at least, increase inequity within countries, particularly affecting the poorest of the poor. This is so because rapid export-oriented growth in developing countries reduces somewhat the differences in income with developed countries, but within countries opportunities vary greatly. Thus, for example, China, which had made enormous gains in reducing hunger, might well suffer an increase in hunger as the income gap between regions increases, as employment opportunities expand in export manufacturing and services but decline in local manufacturing and agriculture, and overall, the safety net system diminishes. Finally, the hunger line shifts relatively as diets expand (by preferences for both animal products and imported products or brands), more and more basic food enters the market, and the income requirement to meet these new needs increases.

Globalization and Reducing Episodic Hunger

Applying similar judgments to episodic hunger (Table 4), the reduction of famine-determined hunger from natural hazards is a great recent success story of a globalized emergency food aid system that relies on both public and private efforts. Today, famine-inspired hunger exists only where war and violent conflict persist. Globalization, however, increases famine vulnerability in the sense of entitlement shifts, as Sen (1981) has shown, especially in cases where the availability of food and purchasing power of rural landless workers can be diminished by far-off events. Globalization has increased the incidence of war and civil conflict both by making weapons easily available (the ubiquitous Kalashnikov) and by diminishing the impacts of war by providing emergency food aid. Recent financial crises often triggered by globalized movements of capital have created sudden episodic hunger in countries where such episodes were rare—as in Southeast Asia, which had made marked progress in reducing hunger prior to the crises. Finally, structural adjustment efforts initiated either internally or at the behest of the IMF almost always lead to an increase in hunger from a decrease in social services and programs, despite some countereffforts.

TABLE 4 Globalization's Impacts on Reducing Episodic Hunger

CAUSAL ELEMENT	HELP	HINDER
Famine	+++	+
War	+	+
Financial Crisis		+++
Structural Adjustment		++

Globalization and Reducing Special-Needs Hunger

Concerning special-needs hunger (Table 5), global efforts to address some major causes of child undernutrition, especially from sickness and disease, by addressing immunization, treatment of diarrhea, and breast-feeding have helped to reduce the rate of wasting and stunting of children (although the actual numbers have increased as a result of population growth centered in the youngest ages). Similarly, the major micronutrient deficiencies of iodine, vitamin A, and iron have been reduced by international programs to encourage iodizing salt, to increase intakes of

TABLE 5 Globalization's Impacts on Reducing Special-Needs Hunger

HUNGER TYPE	HELP	HINDER
Mothers and Children	+++ +++	+ +
Iron, Iodine, Vitamin A		

vitamin A through vitamin A-rich foods and through vitamin A supplementation, and to a much lesser extent, to reduce anemia by iron supplements. In some cases the diversification of diets has helped as well—for example, by providing greater access to iodized salt.

The major countercurrent related to globalization is similar to the previous case: Structural adjustment and diminished development aid have severely constrained many programs directed at addressing these special needs.

PRESERVING LIFE-SUPPORT SYSTEMS

The life-support systems of the planet are often factored into major media: atmosphere, freshwater, oceans, and the biota as biomes, ecosystems, and species. The major threats to atmosphere, freshwater, oceans, and the biota are threefold: (1) the large-scale introduction of pollutants, such as acid rain and chlorofluorocarbons in the atmosphere, heavy metals in the soil, or chemicals in groundwater; (2) the massive assault on biota, such as deforestation in the tropics and the mountains, desertification in dry lands, overfishing of marine resources, and species extinction everywhere; and (3) human-induced climate change.

These threats are incredibly recent. In nine of twelve indicators of global environmental change, half of all the change that took place over the last ten thousand years occurred in our lifetime (Table 6).

DRIVING FORCES OF ENVIRONMENTAL CHANGE

These changes coincide but are not necessarily caused by the most recent wave of globalization.

A general consensus among scientists posits that growth in population, in affluence, and in technology are jointly major driving forces for such change and related environmental problems. This has become enshrined in a useful, albeit overly simplified, identity known as IPAT, first published in 1972 by Ehrlich and Holdren in *Environment magazine* in response to a more limited version by Commoner et al. (1971). In this identity, various forms of environmental or resource impacts (I) equals population (P) times affluence (A), usually income per capita, times the impacts per unit of income as determined by technology (T) and the institutions that use it. Academic debate has now shifted from the greater or lesser importance of each of these driving forces of environmental degradation or resource depletion, to debate about their interaction and the ultimate forces that drive them.

Let me introduce a variant of the IPAT identity (Figure 4)—which might be called the PC version—and restating that identity in terms of population and consumption, it would be I = P*C/P*I/C I equals environmental degradation and/or resource depletion; P equals the number of people or households; and C equals the consumption per person of energy, materials, and information.

TABLE 6 10,000 Years of Environmental Change: Selected Indicators

	SELECTED INDICATORS	
10,000 YEARS OF ENVIRONMENTAL CHANGE	RATE OF CHANGE STILL ACCELERATING	RATE OF CHANGE NOW DECELERATING
Half occurred before our lifetimes	Deforested area Soil area loss	Terrestrial vertebrate diversity
Half occurred during our Lifetimes	Carbon releases Nitrogen releases Floral diversity Sediment flows Water withdrawals	Carbon tetrachloride releases Lead releases Marine mammal diversity Sulphur releases

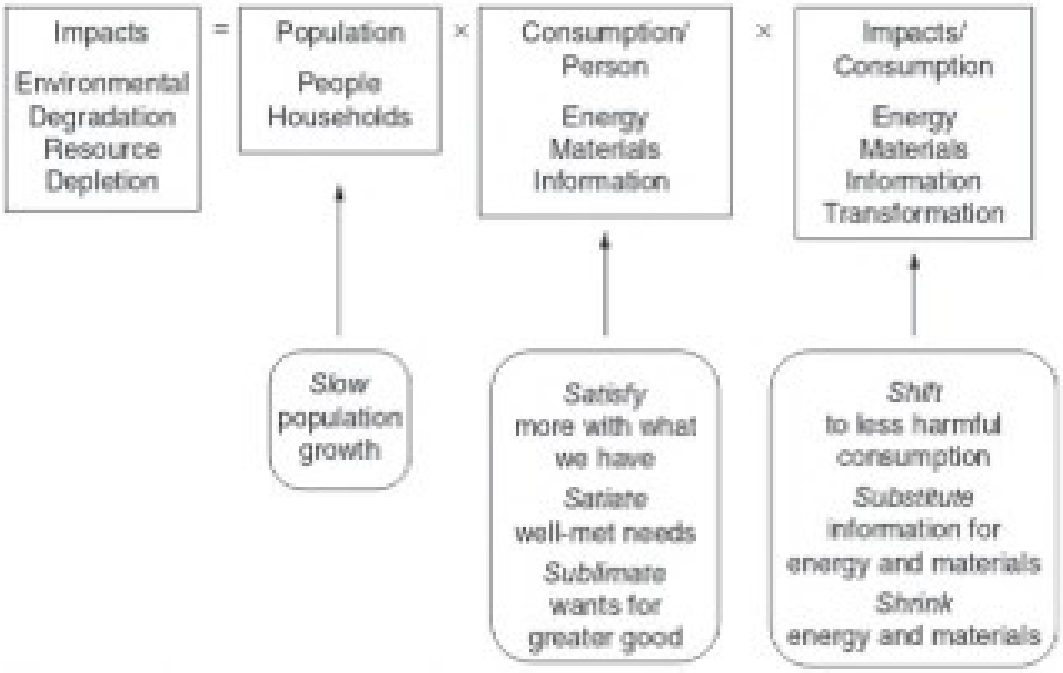

FIGURE 4 Variant of IPAT Identity.

With such an identity as a template and with the goal of reducing environmentally degrading and resource-depleting influences, there are at least seven major directions for research and policy. To reduce the level of impacts per unit of consumption, separate more damaging consumption from less harmful forms and *shift* to these, *shrink* the amounts of environmentally damaging energy and materials per unit of consumption, and *substitute* information for energy and materials. To reduce consumption per person or household, *satisfy* more with what is already had, *satiate* well-met consumption needs, and *sublimate* wants for a greater good. Finally, *slow* population growth and then stabilize population numbers as indicated earlier.

Before using these proximate determinants of the identity, a bit of caution is in order. IPAT is very useful, but is more complex than a simple identity. The PAT terms are only proximate; each in turn is driven by diverse underlying processes. Nor are the PAT terms independent of each other; for example, income (or affluence) influences the rate of population growth and consumption, as well as the technologies used to produce it. Indeed, the supposed technology term is really a catch-all of all the diverse items that determine a different set of impacts per unit of consumption—including technology, but also all kinds of ideas and institutions. Nor do I think the PAT terms are sufficient to examine the impact of globalization on Earth's life-support systems.

To add to the classic determinants, let me borrow from an analysis by Clark (2000) of environmental globalization, one that includes not only the globalization of environmental "stuff"—the energy, materials, biota transformed by production and consumption—but also the globalization of environmental ideas and governance. Specifically, Clark notes three major ideas: planetary management, risk assessment, and sustainable development; along with three forms of governance: by governments; by nongovernmental organizations, both profit and nonprofit; and by coalitions and networks that bring them together.

GLOBALIZATION AND PRESERVING LIFE-SUPPORT SYSTEMS

Thus, there are five causal elements related to globalization's impact on preserving the life-support systems of the planet: population, consumption, technology, ideas, and governance. Examining these five causal elements (Table 7), beginning with *population*, globalization, as

TABLE 7 Globalization's Impacts on Preserving Life-Support Systems

CAUSAL ELEMENT	HELP	HINDER
Population	++	
Consumption		++++
Technology	+++	+
Ideas	+++	++
Governance	++	++

noted, will encourage a decline in the population growth rate, thus lessening future human impacts on the environment. But globalization also accelerates *consumption* in three ways. Insofar as globalization increases income and to a degree inequity, it encourages greater production and consumption. Insofar as it extends the reach of trade and transport, it makes possible consumption of distant resources that would not be consumed locally and the import of goods not available locally. With the spread of ideas and images of "Western material standards of living," it further encourages consumption. Much of this enhanced consumption is desperately needed and desired by the poorer peoples of the world, but much of it will be in the form of material and energy transformations that are environmentally degrading and resource-depleting.

At the same time, globalization facilitates both the creation and the diffusion of *technologies* that lessen the need for energy and materials per unit of production or consumption, that create fewer toxics and pollutants, or that substitute information for energy and material use. Countering these helpful technologies are trickle-down technologies that export secondhand or second-rate technologies, buildings, and toxic or polluting processes to regions heretofore relatively free of degrading or depleting activities.

In the realm of *ideas*, there are also dual impacts. Development is seen as a good to be pursued, with environmental concerns a luxury that developing countries cannot readily afford. Countering this has been the rapid spread of the major environmental ideas of planetary and risk management and sustainable development, to the extent that while rhetoric persists as to differences between the North and the South, international surveys show little difference in the high environmental aspirations among the people of both North and South.

As to *governance*, while there are some 180 international environmental treaties in force to date, they are insufficient to counter the major threats to the atmosphere, oceans, and biota, with the possible exception of the Montreal protocol on ozone depletion. But as Clark notes, governance is greater than governments, and the governance activities of advocacy coalitions and the discourses between transnational corporations and international environmental groups can be very substantial in helping to maintain the planet's life-support system. Indeed, it is often the ability to draw global attention to some local threat that leads to mitigative or preventive actions. All of these forms of governance will grow with globalization. Substantially countering these forms are the notions of free trade unhindered by environmental constraint or regulation currently ensconced in the operation of the World Trade Organization and in various regional trade treaties.

GLOBALIZATION AND TWO-ARMED SCIENTISTS

The Maine version of a well-known story has Senator Edmund Muskie listening to testimony on the need for clean-air legislation. After a series of cautious scientific testimonies by scientists, he asks whether anyone in the room is a "one-armed scientist" who can testify without the endless academic qualifications of "on the one hand ... then on the other hand..."

It is obvious that I must plead guilty to being a two-armed scientist. But it is difficult for anyone thoughtfully confronting globalization not to be. Indeed, whether you sing paeans to globalization as Tom Friedman (2000) does or fear "the manic logic of global capitalism" as Bill Grieder (1997) does, the other hand, whichever it might be, is always present. What may separate the two more is that Friedman thinks that globalization itself will in time deal with its many harms, whereas Grieder

believes in the need for radical revision and strong local to international action.

So, what can the two-armed scientist conclude from this first, tentative attempt to assess the impact of globalization on a sustainability transition: meeting human needs, reducing hunger and poverty, preserving the life-support systems of the planet? In this brief assessment I have focused on how globalization has affected the reduction of hunger and meeting the human needs for food and nurture and on the driving forces of global environmental change.

In sum, globalization to date has not helped reduce the numbers of chronically hungry in the world, although the proportion decreases as population grows. The shifts in investment, income, and job opportunities in some parts of the world are matched by the growth of hunger elsewhere. Episodic and special-needs hunger have benefited more as globalized public and private programs have expanded to respond rapidly to famine, to the special needs of children, and to two of the three major micronutrient diseases. Finally, the increase in hunger from globalized financial crises and policy decisions argues for new sources of instability for the fragile existence of the poorest of the poor. Short-run and long-run simulations of halving hunger argue for these trends to persist.

In preserving the life-support systems, the crucial issue is globalization's impact on current and projected production and consumption of energy and materials that are environmentally degrading and resource depleting. To date, despite major gains in technologies that reduce the use of energy and materials per unit of production, the absolute growth in consumption overwhelms the steady global technological progress. The globalization of environmental ideas has been truly remarkable, and these, along with feminism and human rights, constitute the major ideological revolutions of the contemporary period of globalization. The spread of these ideas has been well facilitated by global interconnections. Similarly, the rise of environmental governance writ large to include corporate behavior and local and international popular initiatives is facilitated by the interconnectedness of globalization and is a great portent for the future. But to date, these have all been insufficient to counter the major threats to the life-support systems, and when projected to address the extraordinary increases in the consumption of the future, such helpful developments may fail as well.

CIVILIZING GLOBALIZATION

If the goods of globalization are to be realized for most of the world's peoples, and if its ills are to be reduced for those people and natural systems most vulnerable, then globalization itself must change in significant ways. Gerry Helleiner (2000), the Canadian economist, as well as Held et al. (1999), call this process "civilizing" globalization. All of them focus on changing governance. Helleiner asks, "Can the global economy be civilized?" "Globalized markets," he argues, "operate within politically defined rules and governance institutions. The current global rules and economic governance institutions are in need of repair, updating and relegitimization," and he goes on to suggest some of those initiatives.

Held et al. (1999) focus not on economy but on politics and three major approaches to civilizing and democratizing contemporary globalization: a liberal-internationalism approach to reform global governance; a cosmopolitan, democratic approach to reconstruct global governance; and a radical republican approach to create alternative structures of governance. Exploring these issues is well beyond the scope of this chapter, but I conclude with my own thoughts on the efforts to "civilize" globalization.

CIVILIZING U.S. CAPITALISM: 100 YEARS

As I contemplate the varied efforts to understand globalization and how to reform or reconstruct it, I am struck by the similarity of contemporary globalization to an earlier period of American history. I perceive an analog to the current global situation in the U.S. economic history of the post-Civil War era. At the time, a truly integrated and nationalized industrial capitalism was created, spurred by the growth of interconnections of

railroad and telegraph. Reading of this age of "robber-baron" capitalism, I am struck by the parallels with current globalization, for example, the great uncertainty at the time, even by the major participants, in understanding the new and different systems that had emerged.

Whatever emerged from the fierce competition and growing monopolization of economic power was initially marred by an absence of law and regulation. Much early regulatory effort was needed around what today are called "rule of law" and "transparency" issues, such as those that make contracts enforceable or stock certificates verifiable. These initial efforts were followed by numerous attempts to control monopoly power and maintain genuine competitiveness.

Almost parallel to these efforts, but much slower, were those that recognized the victims, harms, inequities, and externalities generated by the new integrated industrial system. The initial focus, as it is currently, was on child labor, followed by other working conditions, and then the essentials of a social safety net including disability insurance, unemployment insurance, pensions, and the like. Efforts to gain workers' rights began early and were repeatedly rebuffed, and it was not until 1935 that they were basically recognized. Finally, environmental issues were not really recognized until Earth Day 1970, thus capping a century of effort to civilize U.S. capitalism. Throughout this whole period the individual states of the United States served as the focus for innovative leadership and the testing ground for appropriate regulation that only later was emulated or taken over by the federal government.

The challenge for civilizing globalization is to reduce these century-long efforts to, at most, two decades. The good news is that this effort is well under way, again with a push for "humanizing globalization" or "globalization from the bottom." I illustrate with Bangor, Maine, population 35,000.

CIVILIZING GLOBALIZATION: IN BANGOR, MAINE

Beginning in 1991, Bangor became a sister-city to Carasque, a small village in the highlands of El Salvador, as part of a group of twenty-nine such sister-cities all trying to support returned populations who had been refugees from the civil war. From Bangor, Carasque has received a rebuilt truck, pedal sewing machines, school and health materials, and, when needed, support for human and economic rights in the form of communications expressed to both the U.S. and El Salvador governments. From Carasque, Bangor has received an opportunity to learn of the realities of poor people in the developing world, a model of how young people can exercise leadership in their own communities, and instructive experience as to what alternative schooling might provide.

Trying to find an issue in common to both communities, people in Bangor identified sweatshop—or *maquiladora*—made clothes, the making of which takes jobs from Mainers and exploits Salvadorans. Today, there are a thousand consumers and thirty businesses in Bangor displaying the "clean clothes" sweatshop-free logo, and four Maine communities are now committed to selective sweatshop-free city purchasing. This has led to a statewide selective purchasing act for textiles and footwear that is now being implemented. The "clean clothes" criteria for such selective purchasing—protecting children and workers, providing a living wage, and giving workers rights to bargain—is an agenda for civilizing globalization. Finally, many of the same Mainers recently took a five-hour trip to Quebec City, Canada, for the so-called Americas Summit to protest a free trade area until it incorporates children's, workers', and environmental concerns.

GLOBALIZATION AND A TRANSITION TOWARD SUSTAINABILITY

The very notion of a transition toward sustainability and the concept of sustainable development from which it derives are products of the widening, deepening, and speeding up of the interconnectedness that characterizes globalization. To think of humankind as a whole, to see its links to the fate of the blue planet, has been an essential

part of the globalization process. Thus, those of us who aspire to a transition toward sustainability, to meet human needs while preserving the life support of the planet, cannot be "against globalization." But what we can say (and are saying) is that for globalization to continue to encourage such a transition, it will need to redouble those aspects of this complex movement in our collective lives that help a sustainability transition and to dampen those that hinder such a transition. It is this nexus that will make real the promise of the Neem tree.

REFERENCES

Anon. 2000. "Neem tree free," *The Ecologist*, 30/4:8.

Clark, William C. 2000. "Environmental globalization." In Joseph S. Nye Jr. and John D. Donahue, eds., *Governance in a Globalizing World*. Washington, D.C.: Brookings Press.

Commoner, B., M. Corr, and P. Stamler. 1971. "The causes of pollution," *Environment*, April, pp. 2–19.

Deininger, Klaus, and Lyn Squire. 1996. "A new data set measuring income inequality," *World Bank Economic Review* 10(3):565–591.

Ehrlich, Paul, and John Holdren. 1972. "Review of The Closing Circle," *Environment* 14(3):24–39.

Food and Agriculture Organization of the United Nations (FAO). 2000. *Agriculture: Towards 2015/30*. Technical Interim Report. Rome: FAO.

Friedman, Thomas L. 2000. *The Lexus and the Olive Tree*. New York: Anchor Books.

Grieder, William. 1997. *One World Ready or Not: The Manic Logic of Global Capitalism*. New York: Simon and Schuster.

Held, David, Anthony McGrew, David Goldblatt, and Jonathan Perraton. 1999. *Global Transformations: Politics, Economics, and Culture*. Stanford, Calif: Stanford University Press.

Helleiner, Gerald K. 2000. "Markets, politics and globalization: Can the global economy be civilized?" The Tenth Raul Prebisch Lecture, Geneva, 11 December.

InterAcademy Panel (IAP) on International Issues. 2000. *Transition to Sustainability in the 21st Century*. Conference of the World's Scientific Academies. 15–18 May 2000, Tokyo, Japan, http://interacademies.net/intracad/tokyo2000.nsf/ all/home.

International Council for Science (ICSU) and the World Federation of Engineering Organizations (WFEO). 2002. *Role and Contributions of the Scientific and Technological Community (S&TC) to Sustainable Development*. World Summit on Sustainable Development Secretary-General's Note for the Multi-Stake Holder Dialogue Segment of the Second Preparatory Committee. Addendum No. 8: Dialogue Paper by Scientific and Technological Communities. United Nations Economic and Social Council E/CN.17/2002/PC.2/6. Add.8. Advance Copy, 28 January 2002.

International Monetary Fund (IMF), Organisation for Economic Cooperation and Development (OECD), United Nations (UN), and World Bank Group (WB). 2000. *A Better World for All: Progress towards the International Development Goals*. Washington, D.C.: IMF, OECD, UN, WB.

Kates, Robert W., William C. Clark, Robert Corell, J. Michael Hall, Carlo C. Jaeger, Ian Lowe, James J. McCarthy, Hans Joachim Schellnhuber, Bert Bolin, Nancy M. Dickson, Sylvie Faucheux, Gilberto C. Gallopin, Arnulf Gruebler, Brian Huntley, Jill Jager, Narpat S. Jodha, Roger E. Kasperson, Akin Mabogunje, Pamela Matson, Harold Mooney, Berrien Moore III, Timothy O'Riordan, and Uno Svedin. 2001. "Sustainability science." *Science* 292:641–642, April 27. http://sustainabilityscience.org/keydocs/fulltext/wssd_stc_020128.pdf.

Madison, Angus. 1995. *Monitoring the World Economy 1820–1992*. Paris: Organization for Economic Cooperation and Development.

Mathews, Jessica. 1997. "Power shift," *Foreign Affairs* 76(1):50–66.

National Research Council, Board on Science and Technology for International Development (BOSTID). 1992. *Neem: A Tree for Solving Global Problems*. Washington, D.C.: National Academy Press.

National Research Council, Board on Sustainable Development (NRC-BSD). 1999. *Our Common Journey: A Transition Toward Sustainability*. Washington, D.C.: National Academy Press.

Raskin, Paul, Gilberto Gallopin, Pablo Gutman, Alan Hammond, and Rob Swart. 1998. *Bending the Curve: Toward Global Sustainability, a Report of the Global Scenario Group*. Polestar Series, Report No. 8. Boston: Stockholm Environmental Institute.

Sen, Amartya. 1981. *Poverty and Famines: An Essay on Entitlement and Deprivation*. New York: Oxford University Press.

Turner, B. L. II, William C. Clark, Robert W. Kates, John F. Richards, Jessica T. Mathews, and William B. Meyer, eds. 1990. *The Earth as Transformed by Human Action: Global and Regional Changes in the Biosphere over the Past 300 Years*. Cambridge: Cambridge University Press.

World Commission on Environment and Development (WCED). 1987. *Our Common Future*. New York: Oxford University Press. (Brundtland Report).

2.6.2 GLOBAL WARMING

DONALD AITKEN

There is general agreement that the earth's climate is changing—although the idea that human actions have an impact on that process is now hotly debated in the political sphere, if not among most scientists. This selection considers the measurable changes to our planet from a "Gaian" perspective, which attempts to understand how earth systems respond to normal and unnatural shifts.

The fact that the world's population is increasing, along with current consumption patterns and concentrations of human habitation, makes many believe that we are rapidly approaching a "tipping point," after which it might not be possible to sustain future growth or recover from extraordinary climatic changes.

Drawing on his background in physics with a long-term focus on alternative energy and sustainable architecture, Donald Aitken examines how we generate and consume energy as well as our various attempts to mitigate potentially its dangerous effects.

Since the last ice age, the earth's average temperature over the past 10,000 years has been relatively stable, but many scientists are concerned that the upward trend in average temperatures over the past few decades is likely to continue and cause serious problems for the planet. Other indicators such as higher-intensity storms, floods, and droughts also suggest potentially dangerous effects of climate change. Although there are certainly costs associated with attempting to mitigate the effects of climate change, Aitken argues that the alternative—business as usual—would be far more costly and potentially devastating.

FOCUS QUESTIONS

1. List and discuss several indicators mentioned in the selection suggest that the climate changes occurring today are unusual and are likely the result of human actions.
2. In terms of energy consumption in the past and present, how does the author propose we take collective action based on observable long-term weather patterns?
3. A number of international efforts related to climate change and global warming are reviewed in this selection. Describe two of them and discuss their specific objectives, the extent to which they were supported by different regions of the world, and their perceived effectiveness.

Source: Crist, Eileen, and H. Bruce Rinker, eds. foreword by Bill McKibben., *Gaia in Turmoil: Climate Change, Biodepletion, and Earth Ethics in an Age of Crisis*, pp. 125–150, © 2009 Massachusetts Institute of Technology, by permission of The MIT Press.

KEYWORDS

acidification, albedo, biomass, carbon neutral, energy wedge, Gaian system, global warming, greenhouse gases, hydropower, ice-melt rate, IPCC, Kyoto Protocol, life zone, Montreal Protocol, PV, subsidies (direct and hidden), sweet spot, tipping point

The Gaian system refers to the interconnected natural responses of the Earth to restabilize living and physical systems when perturbed beyond normal bounds. Gaia is, of course, not limited to providing only for human beings, but humans have a huge stake in the outcomes. Nowhere is the strange abandonment of the Gaia stabilization responsibilities of humans more evident or consequential than in our use of energy, in particular in the burning of fossil fuels and its resulting climate destabilization. While many other environmental problems may be deemed of equal or greater urgency, the interaction of anthropogenic climate destabilizations with all natural and human systems is leading many scientists today to identify climate change (and its cause, global warming) as the most pressing environmental issue that needs to be addressed by global cooperation.

THE PIVOTAL ENERGY ROLE OF THE EARTH'S ATMOSPHERE

This article explores the nature of global warming and some of the impacts that are becoming evident today and appear to be heading the Earth system toward *tipping points* beyond which recovery by human actions will not be possible. I review the presently inadequate international response. The scientifically agreed-upon upper limits for the increase in global temperature and the concurrent maximum concentration of carbon dioxide in the atmosphere are presented to underscore the need for the adoption of stringent new policy and energy transition timetables for all nations. I conclude with a brief overview of energy solutions, focusing on the enormous potential of the same renewable energies that have been utilized in the Earth system during its entire history.

All of the solar energy that is absorbed, used, and reused by the Earth's physical and living systems is ultimately re-radiated out to space, only having resided for a time within the Earth's cycles and masses, including a tiny fraction parked in the Earth's vegetation and other life forms. If this outgoing radiation did not exactly equal the incoming energy from the sun, on average, we would either be a frozen, presently lifeless planet like Mars, or we would be an unbearable oven like Venus. After all, all three planets are within the "life zone" around the sun, the region in the solar system in which conditions for the emergence of life could arise. It is the physical properties of the Earth's atmosphere, then, that have been pivotal to making the Earth's temperature and climates livable and suitable for the development of life forms—including us. Tamper with those properties and we are tampering with all life-supporting systems on Earth.

Carbon dioxide along with the other greenhouse gases (e.g., water vapor, nitrogen oxide, methane, and anthropogenic chlorinated compounds) play a major role in the regulation of the flow of energy through the atmosphere from the sun to the Earth, and the counterflow of reradiated energy from the Earth back out to space. The flows must remain in balance to maintain thermal equilibrium. Even though these flows are individually substantial, if they are just slightly mismatched, the destabilizing effects can be great, for an altered energy balance requires that the entire energy equilibrium of the Earth and its physical and living systems must change.

The flows do get unbalanced from time to time. The sun goes through small oscillations in brightness. The Earth's axis periodically changes its relationship to its orbit over long periods. And volcanoes erupt, injecting huge amounts of both dust and carbon dioxide into the atmosphere. The 1991 eruption of Mt. Pinatubo in the Philippines, for example, injected between 15 and 30 million tons of sulfur dioxide gas into the air. In two weeks the cloud had gone around the Earth, and a two-year global cooling was

launched. Volcanic eruptions come and go, as do other surface events on Earth. Over time they average out. But what humans are now doing with the burning of fossil fuels is not averaging out; carbon dioxide is accumulating in the Earth's atmosphere. From the standpoint of the Earth system, this has gone out of bounds leading to increasing destabilizations of the planet's energy, temperature, and climate systems.

FOSSIL-FUEL BURNING AND THE ATMOSPHERIC CARBON DIOXIDE BALANCE

On Mars virtually all of the carbon dioxide is trapped in its soils and rocks, and it is very cold with an average temperature of –50° Celsius. On Venus 96 percent of the atmosphere is carbon dioxide, creating a thermal blanket that surrounds and bakes that planet at an average temperature of +420°C. On Earth, on the other hand, only about 0.04 percent of the atmosphere is carbon dioxide, but even this minute amount plays a critical role in enabling our planet to stabilize at an ambient average temperature of about 14.6°C (58° Fahrenheit).

Until recently we did not know how sensitive the Earth's temperature and climates were to this small but evidently critical amount of carbon dioxide; we are now finding that out as we pour billions of tons of carbon into the atmosphere. When fossil fuels are burned, the carbon that has long dwelled in them is released as carbon dioxide into the atmosphere. This is now yielding an average of 6.1 billion tons of new carbon into the Earth's cycles (about 1 metric ton of carbon for every person on Earth). Even though much of this is absorbed into the oceans, and some is also absorbed by a general increase in the growth rate of photosynthetic plant life on the Earth's surface, about 3.5 billion metric tons of that new carbon is being added to the atmosphere each year. The consistent accumulation of these small amounts of excess CO_2 over the past 150 years, however, has added up to about 39 percent more carbon dioxide in the atmosphere, than would otherwise be there from natural processes alone (IPCC 2007).

Aerosols (dust and small particles) from fossil fuel combustion reduce the flow of the incoming solar radiation while the carbon dioxide product of that combustion, along with the other greenhouse gases injected into the atmosphere by human actions, retard the flows of the outgoing radiation. The two flows are not equally affected, with the outgoing radiation flow impacted the most. The result has been a gradual *net* accumulation of excess energy on the Earth's surface and oceans at the rate of a little under 1 watt/meter squared, averaged over the entire Earth. While seemingly small in absolute amount, this imbalance rate, if it had existed and continued unchecked during the previous 10,000 years, would have raised the Earth's temperature more than 100°C and boiled the oceans.

The amount of carbon dioxide in the atmosphere, and its effect on the Earth's energy balance, are steadily increasing to levels not seen for perhaps a million years or more, and growing at an extraordinarily rapid rate compared with geological history during the development of the human species. Both the present *change* and the *rate of change* of atmospheric carbon dioxide are leading scientists to regard our predicament as ominous.

SWIFTLY UNFOLDING CONSEQUENCES

A number of research paths are converging on the history of atmospheric carbon dioxide content and Earth temperatures that can be traced for up to 750,000 years. The results show a straightforward relationship between carbon dioxide and temperature: as the one increased or decreased, so did the other. While the precise cause-and-effect details are not clear, the overall patterns are clear, and can reasonably be, expected to continue to in the same relationship into the future (Figure 1).

Equally clear from these various research results is that the burning of fossil fuels has taken the Earth's atmosphere, and hence energy flows and balances, into uncharted waters. Furthermore, since this has all happened in a short 150 years or so, the rise in carbon dioxide and temperature show up as sudden spikes at the end of charts of Earth's recent temperature and

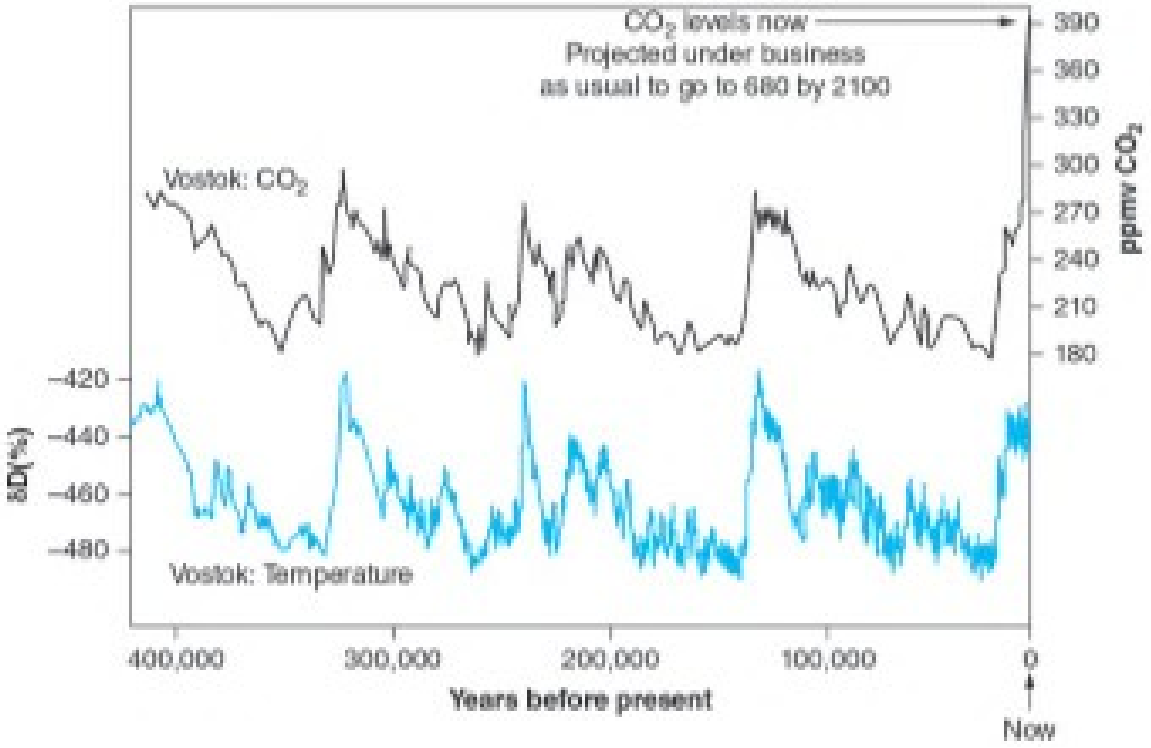

FIGURE 1 Concentrations of CO_2 and Earth temperature based on Vostok ice core samples

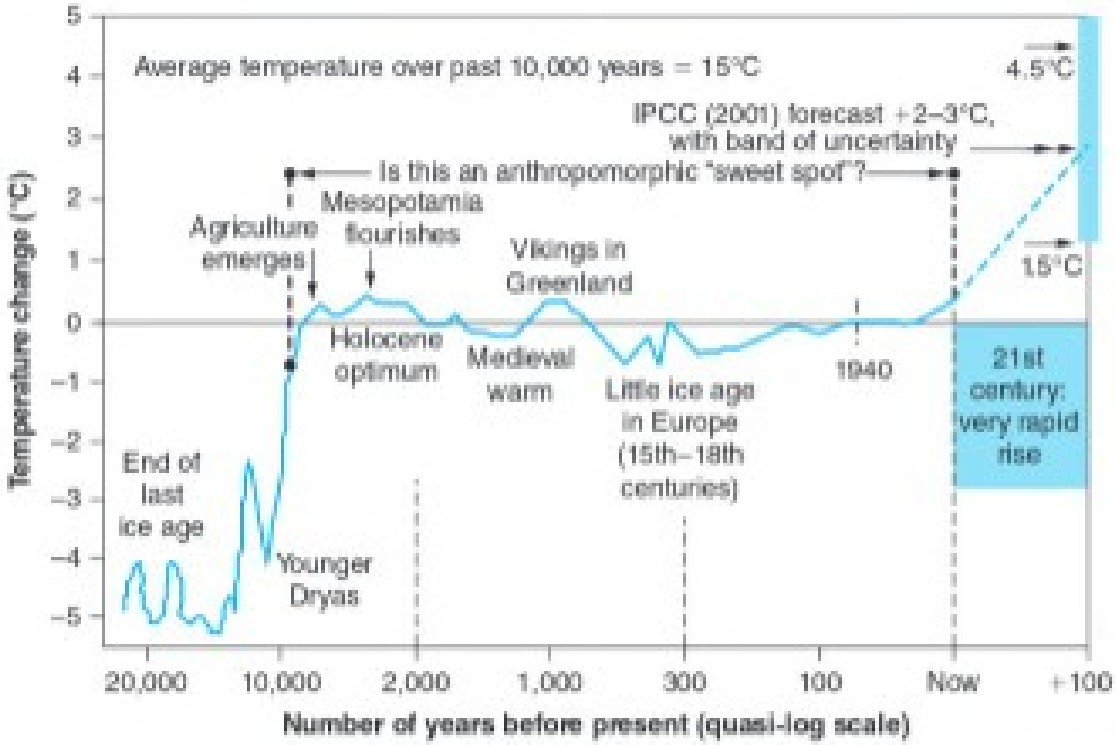

FIGURE 2 The last 10,000 years have been ideal for the development of human societies

CO_2 histories. Both of those spikes are rising rapidly as the world's old and new fossil-fuel power plants (primarily the coal plants), as well as other sources (e.g., transportation) continue to increase levels of carbon dioxide in the atmosphere (Figure 2).

Present and projected future concentrations of carbon dioxide depicted on the 400,000-year Vostok ice core data sets for both Earth temperature and atmospheric carbon dioxide content. The IPPC is the Intergovernmental Panel on Climate Change. Figure courtesy of Dr. Robert Correll.

So what can we expect to be the future results of our actions? Certainly a rapid warming of the Earth now appears underway. But since "climate" is the Earth's mechanism for the redistribution of its surface energies, it is equally inevitable that the Earth's climate must change as well.

How significant might that change be? For the past 10,000 years, following the exit out of the last ice age, the Earth's temperature has been remarkably mild and stable—nicknamed a "sweet spot" by climate scientist Robert Correll—not increasing or decreasing more than 0.5°C (see Lempinen 2007). This set the climatological stage for the evolution of great civilizations. The IPPC analyses, however, suggest that by the end of this century, in the absence of stringent global control of greenhouse gas emissions, the Earth's temperature could climb as much as 4.5°C above that 10,000 year "sweet spot" average. This is roughly equal to the temperature difference between the last ice age and today, demonstrating how only a few degrees of warming or cooling can have extraordinary consequences for the biosphere and for human civilizations within it.

The last 10,000 years seems to have been ideal for the development of human societies. Is this an historic "sweet spot" that enabled humans to flourish? Figure and caption text by Dr. Robert Correll.

We do not need to wait for decades to discover how sensitive the Earth is to the impacts of our fossil fuel burnings, nor how quickly the Earth's energy systems can be unbalanced. The average temperature has only risen by 0.74°C (about 1.3°F) in the past 100 years. What is remarkable is the nature and pace of changes *already* taking place as the result of this small change, and how much more rapidly those changes are happening than even the projections of the best computer models. This would suggest that, while the Earth system is inherently robust, its balances are also very finely tuned.

Rather than offer a litany of all that *might* happen, then, I would rather lean on a few examples of what is happening *already* to underscore the urgency of the needed human response to the imperiling of the Earth systems on which both human beings and our contemporaneous species and ecosystems so vitally depend.

ARE HURRICANES TELLING US SOMETHING?

We remember all too well how Hurricane Katrina devastated New Orleans and other Gulf coastal regions in 2005, bringing untold suffering and damage. The previous year had seen four major hurricanes striking the US shores. Was global warming the cause? While it cannot be proved to be the cause of any particular storm, the accumulating evidence is showing a disturbing overall *pattern of climate change*—storms, floods, droughts—tracking the rate of carbon dioxide increase in the atmosphere.

For example, evidence shows that the number of the *most intense* tropical storms has increased by 80 percent over the past 35 years, and that the average intensity (which produces the damage) of storms created in the Atlantic Ocean has more than doubled during the 1983 to 2005 period (Webster et al. 2005; Kerr 2006). (Hurricane Wilma, on October 19, 2005, was for a while the strongest hurricane ever recorded with sustained winds of 170 miles an hour.) Hurricanes spawned in the North Atlantic further reveal an unambiguous correlation of increasing hurricane strength with human emissions, scientifically distinguishable from other possible natural causes (Mann and Emanual 2006). Recent theoretical models have improved to the point that some of these observations are predictable, although the models produce more complicated correlations in different ocean basins, and the models are not all in agreement. They do not, however, disagree with the statistical observations of hurricane intensity (Emanual et al. 2008). Hurricanes draw their energy largely from the surface waters over which they pass. Since research has shown unequivocally that the temperatures of the upper layers of the world's oceans are tracking the increase in lower atmospheric temperatures, the correlations of increasing tropical storm and hurricane strengths with increasing ocean surface temperatures is not surprising, but expected.

The damage caused by Hurricane Katrina has variously been estimated to be US$200 billion to $300 billion. Estimates for what it would have cost in coastal protection to prevent most of this

damage are in the range of US$2 billion to $3 billion, or about 1 percent of the cost of the resulting damage. The message here is that prevention and mitigation of global warming impacts will be, in the long run, an economic bargain compared to the costs of inaction.

EVENTS IN THE ARCTIC OCEAN

Perhaps even more alarming than the increasing global temperatures and intensities of storms worldwide are the unexpectedly rapid effects of global warming in the Arctic and Antarctic regions. In the Arctic the average temperature has been climbing for a number of decades twice as fast as elsewhere. The temperature of Alaska has been growing 6 to 8 times faster than the rest of the world.[1] The Antarctic peninsula is warming more rapidly than anywhere else on Earth (Bell 2008). Increased precipitation, shorter and warmer winters, and decreases in snow cover are already being documented in Arctic regions. The reduced solar energy reflection (or albedo) from the loss of ice and snow is moreover expected to accelerate the warming trend, generating a feedback loop that may be leading to runaway meltings and other far-reaching regional changes.

Summer sea ice in the Arctic Ocean has been declining over the past several decades, reaching an all time low in 2005. The steady decline was replaced by a plummeting decline in 2007 when the summer Arctic sea ice decreased by 23 percent from the 2005 low. Further examination led to the startling discovery that more than two-thirds of the rapid sea-ice melting is now happening from below, an amount representing 5 times the normal summer loss, caused by the warming of the waters, rather than from above, as usually caused by the summer sun (Perkins 2007). The 2005 melting had already led climate scientists—who have been carefully modeling the potential impacts of global warming—to note that rather than earlier calculations of taking 100 years for the Arctic Ocean to become ice free in the summer, it could happen by 2050. The 2007 melting has now led some to note that summer sea-ice melting may be 20 years ahead of the theoretical projections; they have further revised their estimate of an ice-free Arctic summer to 2030. Climatologists are surprised at how much faster these consequences are unfolding than even the extreme scenarios of their models. Normally cautious scientists are now stating openly that 2007 may prove to have been the "tipping point"—the point at which the melting of summer Arctic sea ice became self-perpetuating (ibid.).

OBSERVATIONS IN GREENLAND AND THE ANTARCTIC

The melting of floating sea ice alone does not raise sea levels. It is the melting of land-based glaciers, especially those covering Greenland and the Antarctic continent, that will contribute directly to sea-level increase. The Greenland ice sheet would raise the oceans by 24 feet if it melts; a melted west Antarctic ice sheet would increase sea level another 19 feet; and if the east Antarctic ice sheets were to melt sea levels would increase another 170 feet. These calculations give a total of 213 feet potential sea-level rise from melting ice on land. It has long been assumed that these meltings would take millennia.

But here again the effects appear to be accelerating at a pace well beyond those expected from computer modeling. The ice-melt rate in Greenland, for example, from 2004 to 2006 was *250 percent greater* than the ice-melt rate from 2002 to 2004. The record high Greenland ice sheet melt of the last 50 years, set in 2005, was trumped in 2007. Studies have demonstrated that, whereas the melting had been more related to regional climate changes between 1960 and 1990, the pattern of melting has since changed to reflect global temperature variations, putting the fingerprint of global warming firmly on Greenland's ice sheet losses (Hanna et al. 2008).

The disquieting increase in Greenland's ice melt is being matched on the west Antarctic continent. Scientific studies have shown that between 1996 and 2006 there was a 59 to 75 percent increase in annual ice loss from west Antarctica, accompanied by a 140 percent increase in ice losses from the Antarctic peninsula; the ice loss from Antarctica now nearly equals that from

Greenland. The most startling event was the complete collapse of the Rhode Island sized Larsen B iceshelf in a span of two months in 2002, following 10,000 years of stability. The floating ice shelves in front of glaciers and ice streams serve as dams to slow the ice stream motion. As they break up, the ice stream flow into the sea accelerates. In addition summer melt water (appearing more and more on the surface of glaciers) is apparently being conducted to the glacier bases, helping grease, and hence accelerate, the glacier flow rate (Bell 2008). The more rapid flow of glaciers is interpreted as the source of the accelerated ice loss in Antarctica.[2]

Whereas some scientists (see IPPC 2007) still conservatively project less than a 1 meter sea-level rise by the end of this century, the evidence of the rapid and accelerating changes now taking place has led other leading climatologists to suggest that it could be 2 meters or more—with a devastating potential flooding and storm erosion impact for many continental coastal lands, and a death knell for low-lying islands.

THE OCEANS

Global warming has introduced two additional disturbing changes to the world's oceans. First, there is already a measured reduction in primary ocean productivity in nine of the twelve ocean basins. As life itself is a crucial participant in the Earth system, such a destabilization of life in the world's oceans can certainly be expected to ripple through all global systems. This will be particularly difficult for the future human populations that depend on fish protein. (Much is now drawn from aquaculture, but that is also introducing problems with chemicals introduced into fish feed and illnesses in fish raised in captivity.)

Second, drops in oxygen levels, observed throughout the Pacific Ocean and with oxygen-deprived "dead zones" appearing with ever-increasing frequency and intensity near shorelines, are being linked not just to the usual source of overfertilization from agricultural runoff (obviously directly human caused) but apparently also both to the warming of the waters and to changes in ocean circulations that are likely a result of global warming (and hence indirectly caused by human interventions in the atmosphere; Juncosa 2008).

Things can only get worse as, third, the increase in carbon being absorbed by the oceans is already leading to their acidification. At the minimum this will reduce the oceans' ability—as the major global CO_2 sink—to absorb the increasing CO_2. But of greater ecological consequence is that when that acidity reaches a critical level, phytoplankton can begin to calcify and disintegrate. Recent research suggests that primary production could increase from this, while others project that with increased calcification, phytoplankton disintegration would reduce primary productivity. Projections suggest that this could happen with atmospheric CO_2 concentrations in the range of 580 to 720 ppm. Since the World Bank has projected a CO_2 concentration of 750 ppm by 2100 (a level that analysts concur *must* be avoided), this could be a genuinely realistic scenario with potentially disastrous potential consequences for all of the ocean food chains, and thus obviously for people as well.

FIRST INTERNATIONAL RESPONSES

International efforts have been underway to address global warming and climate change, but the results to date have been minor and continue to be mired in politics, leaving little room for the kind of response needed. But there is hope, and emerging evidence that an international response is possible.

The world's governments banded together to address a global atmospheric problem in March 1985 with the adoption of the Vienna Convention for the Protection of the Ozone Layer. This was followed up in September 1987, when twenty-four nations signed the "Montreal Protocol on Substances that Deplete the Ozone Layer," putting into action the aims of the 1985 Vienna Convention by setting legally binding controls and targets for phasing out the chlorinated chemicals that were causing the problem. The results have proved to be amazingly effective. International cooperation avoided a global atmospheric hazard with serious health consequences to humans and other species.

Less than three months after the implementation of the Montreal Protocol to protect the ozone layer, in December 1987 the United Nations General Assembly passed a resolution entitled "Environmental Perspective to the Year 2000 and Beyond," in which global warming first surfaced as an important policy area that should command the international attention.[3] Launched by UN resolution a year later, the "Intergovernmental Panel on Climate Change" (IPCC) was created jointly by the World Meteorological Organization and the United Nations Environment Programme (UNEP). The responsibility of this international assemblage of scientists was to assess the science and risks of "human-induced climate change," and to provide counsel to the United Nations, policy makers, and the public.

The first IPCC assessment report, published in 1990, was crucial in representing an international scientific convergence on the view that human activities are the main driving force behind global warming. That report, in turn, provided the empirical basis for the creation of the next policy framework by the United Nations—the United Nations Framework Convention on Climate Change (UNFCCC). In March 1994 the UNFCCC went into effect with 189 nations as signatories to this Treaty—including the United States. By their signatures, the countries that signed the UNFCCC declared that they were "determined to protect the climate system for present and future generations."[4] There have been three additional IPCC assessment reports since then, the most recent released in the summer of 2007. In addition to steadily refining the scientific conclusions and formalizing recommendations for international response, over the years the four assessment reports have been notable for their increasing certainty that human actions are now the primary cause of global warming and their growing concern regarding both the present and future probable consequences.

In December 1997 the famous Kyoto Protocol was adopted, delineating binding targets and a timetable for greenhouse gas emission opened for signature in March 1998, but not ratified until early 2005, when 55 nations (accounting for 55 percent of the global carbon dioxide emissions) had signed. The goal of the Kyoto Protocol is to reduce the greenhouse gas emissions (in carbon dioxide equivalents) by the developed nations to an overall average of 5.2 percent below the 1990 levels, accomplishing this in the 2008 to 2012 window. Different signatory nations were assigned different targets, with the total to meet the overall goal. The United States refused to sign. Russian compliance put the total signatories over the top, enabling the Kyoto Protocol to become international policy.

The developing nations were given a pass in this agreement, a source of great US displeasure, to prevent economic hardship during development. But with China surpassing the United States in carbon dioxide emissions in 2008, it is clear that a next, post–Kyoto Protocol will need to differentiate between those developing nations (China, but also India) that are now on a par with developed nations in emissions and those that are not.

The thirteenth "Conference of the Parties" (COP13) in 2007 in Bali, which was to begin to develop standards for the post-2012 period, dissolved into political wrangling and rhetoric, with the United States continuing to stand in strong opposition to any mandatory targets. As this is being written, the hopes are being pinned on COP14, to take place in 2009, when, according to agreement by all nations, a threshold is to be firmly agreed upon regarding maximum global atmospheric concentration of greenhouse gases.

The Kyoto Protocol represents a beginning, standing as an important international acknowldgement of the need for cooperative action. However, even if all nations were to meet the Kyoto CO_2 reduction targets, emissions would *continue to grow*. Clearly, more stringent goals and reduction targets for greenhouse gas emissions need to be set and implemented. The question is how? Will it be technically and economically feasible to meet stringent reductions in greenhouse gas emissions in ways that can also facilitate the return to a Gaian balance of Earth

SETTING GLOBAL WARMING AND GREENHOUSE GAS TARGETS FOR THE EARTH SYSTEM—AS WE KNOW IT

According to international scientific consensus, in order to avoid "dangerous climate change," the long-term temperature rise of the planet's surface lands, atmosphere, and waters should not exceed 2°C. We are almost 40 percent of the way to that point already. The maximum global temperature goal that scientists are converging on requires the concentration of CO_2 be no more than 450 ppm by 2050 (Meinshausen 2006). Since we are now at 390 ppm of CO_2—a rise in concentration by about 110 ppm from the pre-industrial levels—this target constitutes a formidable challenge to all nations.

The approximate figure of 450 ppm represents the concentration range of greenhouse gases at which there is a 50 percent chance that the temperature will not exceed the 2°C limit, and a nearly 70 percent chance it will not exceed a 3°C rise. It is the best available estimate of the upper limit for avoiding disastrous ice melting in the Arctic and Antarctic regions, unacceptable ocean level rise, and dangerous climate and ecological changes.[6] In other words, this is a goal for reestablishing a balance of the Earth's physical and living systems within acceptable bounds, and within a range that can again be stabilized and maintained, by Earth system processes.

While there is no certainty that even this figure is safe, it sets a defensible target for action. The European Union has already adopted these figures—2°C and 450 ppm—as the basis for the EU emission reduction goals. These targets indicate that by 2050 global greenhouse gas emissions must be reduced by 50 percent over the base year 2000 levels. But in recognizing that the developing nations will have difficulty in meeting these reduction goals even if they try, the industrial nations—including a cooperating United States—will need to reduce their emissions by an average of 70 to 80 percent below year 2000 levels by 2050.[7]

In order to have a reasonable chance of accomplishing these deep cuts, industrial nations must have their emissions peak by 2010, and thereafter begin to decline at an average rate of 4 percent per year—that is, at a faster rate than the global emissions are presently increasing. If the emissions do not begin to decline until 2020, the decline rate will have to be an almost unachievable 8 percent per year. The inevitable conclusion is that a meaningful response by all nations must be set in motion immediately, if there is to be hope of success. While the path will not be easy, these aggressive targets can be met. A wide range of technology and policy options will be required, woven into international agreements that include significant support by the industrial nations for the developing nations. The starting point, however, must be the international acceptance of CO_2 concentration and temperature goals as binding global limits.

Governments opposing mandatory rules argue that voluntary goals by industries and governments should be adequate. However a national survey of 500 big industries in Britain, the United States, Germany, Japan, India, and China reported in early 2008 that climate change ranked only eighth in the concerns of the business leaders, behind, for example, increasing sales and securing growth in emerging markets—and would probably become of even lesser concern if climate change causes the global economy to deteriorate (David et al. 2008).

Some businesses have become environmentally responsible on their own initiative, recognizing that tackling energy efficiency and greenhouse emission reductions will place them in an advantageous position in the future, at a time when other businesses will have to scramble to meet inevitable new emission rules, including the payment of carbon taxes.[8] But, by a large margin, the businesses polled stated that setting goals for emission reductions was properly the role of governments. The unfortunate present result of that nonperforming "proper role" has been a nearly 2 to 3.5 percent annual *increase* in global CO_2 emissions with no 'immediate prospects of rates abetting. Voluntary controls seem fraught with competing interests and political agendas, and are not adequate in confronting the enormous technical and policy challenge that we face.

The nations that signed onto Kyoto now have the 2008 to 2012 period to deliver their committed emission savings, so this rate of global increase may indeed reduce for a time. On the other hand, the United States and China, neither signatories to the Protocol, and both staunch opponents of mandatory emission reduction standards, together account for nearly 50 percent of global greenhouse emissions. China is bringing online a new coal-fired power plant every one or two weeks, while the United States has over 150 new coal plants in various stages of planning or construction.

If we move our focus away from national government and corporate initiatives to voluntary actions by *local* governments, the story changes. In the United States, for example, on February 16, 2005, the day the Kyoto Accord went into effect, the mayor of Seattle, Washington, launched the US Mayors' Climate Protection Agreement. In June of that year, the US Conference of Mayors passed that agreement unanimously. Their target is modest, fashioned after the then vice president Al Gore's acceptance in Kyoto of a US target of a 7 percent reduction in greenhouse gas emissions from 1990 to 2012. While the United States did not sign the final agreement, by the end of January 2008, 780 cities in all 50 US states and Puerto Rico (representing more than one-third of the US population) had pledged to meet or exceed the terms of the unsigned US Kyoto Protocol obligation.[9]

This precedent was followed up at the December 2007 UN Climate Change Conference (COP13) in Bali. Whereas the participating national governments had great difficulty in securing significant agreements amid much squabbling and posturing, the World Mayors' and Local Governments' Climate Protection Agreement was launched at the same meeting on December 12.[10] A Gaia-appropriate target of reducing global emissions by 60 percent from 1990 levels by 2050, and by 80 percent for the industrial nations, was adopted. Signatories also agreed to offer annual reports on their greenhouse gas emissions

EMISSION-REDUCING STRATEGIES: CONSIDERING THE ALTERNATIVES

Energy efficiency in buildings provides the technically easiest, least expensive, and fastest measure for reducing global warming emissions. Both of the Mayors' Climate Protection Agreements include energy efficiency improvements to city facilities. The US Agreement goes further, in adopting the goals of the "Architecture 2030 Challenge," whereby the fossil fuel reduction standard for all new buildings is to be increased to 60 percent in 2010 all the way to 90 percent by 2025. The aim is to become carbon neutral by 2030.[11] This standard was adopted as policy by the American Institute of Architects and was included in the US Energy Independence and Security Act of 2007 (signed into law at the end of 2007) for all federal buildings.

Buildings in the United States are directly and indirectly responsible for 48 percent of the US greenhouse gas emissions,[12] and use 67 percent of all of the country's electricity. By 2035, however, the amount of new and refurbished buildings in the United States will approximately equal all of the buildings currently in place in the entire country.[13] A nationwide adoption of the Architecture 2030 standards has the potential therefore to displace up to 36 percent of the nation's greenhouse gas emissions by 2030. This efficiency measure alone could account for almost half of the 2050 target of an 80 percent reduction in US emissions.

Generally, the most impacting buildings are in or close to cities. About 50 percent of the world's population lives in cities, today also the locus of 75 percent of the worldwide consumption of energy and 70 percent of the world's consumption of electricity. By 2030, two-thirds of the globe's population is expected to live in cities, accounting for an even larger share of global energy use and emissions. National policy for the reduction of greenhouse gas emissions must consequently begin in the cities. That's where the people are—as well as the buildings, the vehicles, and the opportunities for the most environmentally and economically advantageous

together to bring about change, as in the United States or Spain, could *de facto* begin to define national policy.

Over the next 50 years, a multitude of solutions to reduce global warming emissions will need to be pursued *simultaneously*, each involving serious and concerted effort. Some proposals are pie-in-the-sky, relying on grand unknown new technologies, or proposing giant sunshades in space, or other risky geoengineering with substances injected into the atmosphere or oceans. Others, on the surface, appear to many decision makers to have some merit, such as today's rising enthusiasm for the nuclear power option. There needs to be a way for decision makers to make early assessments of the reality and potential value of technology options that are going to be funded.

More realistic options, based on available technology, and appropriate for the first fifty years of this century's global energy transition, were outlined in 2004 in a well-known paper by Stephen Pacala and Robert Socolow.[14] They derived a series of wedge-shaped pieces of a carbon-emission pie chart, with each wedge worth a reduction in global carbon emissions of 25 billion tons by 2056. Seven such wedges taken together (representing 7 policy and technology tracks) could stabilize global CO_2 emissions at about today's level.[15] This in turn would lead to an ultimate CO_2 level about twice the pre-industrial level (around 560 ppm). On present-day scientific assessments, this is still a dangerously high CO_2 level. The primary value of Pacala and Socolow's work, however, was to convert a set of numerical emission-reduction targets into understandable and feasible numerical requirements for the application of related technology options.

The value of this simple approach can be illustrated by returning to the nuclear power example. Just one wedge would be provided by a doubling of today's nuclear power capacity, and then only if it offsets coal-fired generation. With 435 nuclear power plants in operation worldwide today, and because almost all of those plants will have been retired and decommissioned power plants would be required to become operational every year for the next 50 years. Given the realities of the availability of materials, components, skilled labor, finances, and regulatory oversight, this may represent 3 to 5 times the maximum realistic rate for nuclear plant construction.[16] An alternative way of looking at this is that, within the present maximum rate for new nuclear power plant construction, by 2030 only 1/7 of a wedge will have been built, leaving about 98 percent of the remaining requirement for the reduction of carbon emissions to other policies and technologies.[17]

These simple calculations reveal that in all likelihood, nuclear energy could only provide a fraction of the required carbon-free energy by 2050, probably no more than a few percent. In today's seeming rush to revive nuclear energy, this should provide a note of realism to decision makers who will need to see to it that the other necessary technical options are also adequately pursued (Hultman et al. 2007).

Indeed, when one steps back and looks at future energy resource options, the naturally available renewable energy resources stand out, not only because of their enormous resource potential but also because they are the same ones that have powered Gaia processes for all Earth's history, solar, wind, ocean thermal, biofuel, hydroelectric, geothermal, and tidal. They are moreover the only energy sources that are guaranteed to be available in perpetuity.

Consider the following comparisons: The energy value of all sources stored in the Earth[18] is estimated to be about 9,100,000 terawatt hours.[19] The annual resource potential of renewable energy is about 350,400,000 terawatt hours per year—almost 40 times the total of stored energy in the planet. The annual energy from solar radiation accounts for 99.9 percent of this. Wind energy, while small compared to solar, is still about three times the total of all energy used by humans (which is about 80,000 terawatt hours per year). These figures invite the conclusion that about 0.075 percent of the land area on the planet receives energy equivalent to the total annual use of energy by all human civilizations.

all other long-range resource potentials—and it is renewed every year.

That's the good news. What is often considered to be the bad news is that the renewable energy resources are diffuse. You cannot pour a concentrated amount of any of them out of a can, or make a lot of money by drilling into a solid or liquid resource of it. To gather them requires constructing collection devices over large areas and paying for those collection devices up front, although after that initial investment the energy is free.

Renewable energy systems are still early in their global applications, so that they are more costly than the conventional resources. Price comparisons are highly distorted, however, by both direct and hidden subsidies for the conventional energy sources. Hidden subsidies include the costs of a military presence and wars to protect access to oil, the failure to include the costs of the environmental and human health impacts arising from their use, and the failure to adjust current prices for future risk in supply. The greatest of the latter risks will be the (apparently near-term) beginning of the period when global supply of fossil fuels cannot keep up with growing demand—followed by an absolute decline in supply (peak oil). Other risks include the effect of growing global nationalism and militarism on global energy supplies, accompanied by risks from terrorism to energy infrastructure and nuclear power plants, as well as risks of nuclear terrorism, using materials meant for the production of nuclear power.

Renewable energy applications are blessedly immune from these impacts and risks, providing greater security, human health and energy independence *advantages, and lower impact on natural systems. Renewable energy creates real employment and converts fossil-fuel energy* dollars that previously went out of the cities or the region into local and regional benefits with real economic advantages (also not included in fuel price comparisons). The advantage of fossil fuels, in terms of ease of use and extensive infrastructure, is short term and evanescing. The enormity and benefits of the solar resource, in particular, compared with the finite amounts and serious shortcomings of fossil resources, assures that renewable energy will be the long-range energy resource for human civilizations. Following the transition (to unfold in this century) the climate system will hopefully stabilize at familiar equilibriums—friendly to human life and extant species and ecosystems.

RENEWABLE ENERGY RISING

Something is happening with renewable energy worldwide: the use of those resources as carbon-free technology and policy solution to global warming is growing with astounding rapidity. Solar photovoltaics (PV) and wind generation are now the fastest growing energy sources in the world. Renewable energy supplies over 17 percent of the world's primary energy, if one includes traditional biomass and large hydropower.[20] In 2007 the "new" renewables (i.e., solar, wind-, geothermal, small hydro, modern biomass, and biofuels) accounted for only 2.2 percent of the world demand for primary energy and 3.5 percent of global electricity production. Yet these small numbers mask the dramatic rates of growth for renewable energy capacities. The small percentages of the global total for these renewables also masks the greatly accelerating pace of annual investments in those technologies. In addition the renewable energy industries have been shown to create a proportionately larger number of employment opportunities for unit energy output than the conventional energy sources, in essence turning money previously paid for fuels and large central power plants into money paid for people.[21] The global stock markets have responded favorably to all this, with stock values for the publicly traded renewable energy companies often leading the way in annual growth rates and frequently growing in market value faster than the market as a whole.

The most concrete sign that renewable energy has come of age in the global transition to clean, noncarbon resources is the scope of the national standards now being set for the progressive adoption of renewable energy into the total energy mix. These are not in the 2 percent range, but rather they average in the 10 to 20 percent

range. This policy growth has spurred and supported the market growth of renewables. While a complete overview of renewable energy policy developments is beyond the scope of this chapter, a few highlights are worth noting.

A recent survey[22] has revealed that, by 2007, 61 countries had set national renewable energy supply targets and mandatory timetables for achieving them. In the United States, renewable portfolio standards (RPS), defining the share of renewable energy in the electricity production mix required by certain dates, have been adopted in 50 percent of the states and the District of Columbia, ranging, for example, from a low of 8 percent by 2020 for Pennsylvania to a high of 20 percent by 2010 for California.[23] The European Commission adopted binding goals that extended their 2010 policies to new targets of 20 percent of final energy to come from renewable resources, along with 10 percent of transportation fuels, by 2020. Also included in the list of countries adopting binding renewable energy standards are 13 of the developing countries. China, among them, adopted their plan in September 2007, calling for 15 percent of primary energy from renewable resources by 2020.[24] In addition at least 57 countries, including 21 developing nations, have set binding targets for the share of electrical power to come from renewable energy.

The rise of renewables—from technological innovations to policy shifts—demonstrates that present and long-range directions for the development and implementation of renewable energy are now on a par with major conventional energy resources. Renewable energy has reached adolescence. It is also clear that renewable energy policies, as they emerge from this adolescence, require cooperative parenting skills at every level of government, from cities to states to nations. It is equally clear that society will need to transfer the enormously distorting and wasteful present funding from pathological applications like wars to the productive development and deployment of the energy technologies that will power the future human societies and industries. The funds are there, and available, if channeled by appropriately maturing societal priorities.

CAN RENEWABLE ENERGY MEET THE NEEDS OF THE FUTURE?

Earlier in this article I examined the implications of accomplishing one wedge (25 billion tons of CO_2 displaced over 50 years) from nuclear power. While technically possible, it is probable that the realities of such a necessary pace of construction of new nuclear power plants would make that a difficult target. On the other hand, the nuclear industry feels that it is quite achievable. If it is, achieving this target would still leave 86 percent of the carbon free energy to come from other sources 50 years from now.

What would it take to deliver one wedge of wind energy or solar energy? The authors of the wedge formalism suggest that to displace one wedge of coal-fired generation by either wind or solar technology would require 2,100 GW (2.1 TW) of installed capacity, including sufficient conventional-fired backup to level the intermittent outputs of these two resources. This is a huge amount, representing 23 times the total global installation, of wind power, and 210 times the global installation of solar power, at the end of 2007. How do these fare with the same kinds of realism tests applied earlier to nuclear power?

An analysis was undertaken in 2005 by Stanford University scientists to quantify the world's practical and realistic wind power potential, based on minimum necessary wind velocities for economic energy production and average turbine hub heights. Their research revealed 72 TW of potential (Archer and Jacobson 2005). Tapping into just 3 percent of this would produce one wedge of CO_2 reduction, so the wind resource is there in abundance. The present annual growth rate of wind energy (a doubling every three years) yields that this 3 percent could be achieved in 14 years. While this leads to unrealistic rates of installations, spreading it over 50 years could well be achievable.

As for solar energy for all practical purposes it is unlimited. The one wedge target for solar power systems could be reached in 16 years with a doubling time of half of what occurred during the 2002 to 2006 period. Again, this could become a realistic target if spread over 50 years.

Adding another wedge produced from biomass used directly for heat and electricity, and to produce carbon-free fuels, would allow the world to meet over 40 percent of global energy needs from renewable energy resources after 50 years.

This is getting close to the 2050 global target of a 50 to 60 percent reduction in global CO_2 emissions to avoid the "dangerous climate change" threshold, although it is still well below the necessary 70 to 80 percent reduction target of the developed nations. But as earlier noted, another 40 percent reduction in CO_2 emissions could be accomplished in the developed nations from energy efficiency improvements, in particular with buildings.

One conclusion is that while adding the one wedge of nuclear power would further reduce the CO_2 emissions, that may not be necessary, for the combination of the renewable energy resources and energy efficiency may well be enough *to* avoid the dangerous climate change thresholds. This conclusion is supported by a number of published models and scenarios that appear to show a feasible renewable energy future meeting from 20 to 50 percent of global energy needs in 30 to 50 years, and up to 100 percent of global energy needs by the end of this century (Martinot 2008). For example, a 2007 analysis by scientists of the National Renewable Energy Laboratory in the United States concluded that 57 percent of the necessary reduction in year 2030 CO_2 emissions (for the United States) could be achieved by energy efficiency, with renewable energy technologies providing 50 percent of the residual energy required in that year.[25] These together would produce the necessary reduction in US CO_2 emissions by 2030, continuing to lower after that.[26]

The question, of course, is whether this will all be fast enough to prevent serious climate changes and enormously expensive consequences for human civilization. Recent publications have been sounding the alarm that even these, kinds of legislative actions may no longer be adequate to prevent the most serious consequences of global warming, and that *leaps* in policies for much higher and more aggressive energy efficiency and conversion goals to carbon-free energies are now needed by *all* nations[27] (Hansen 2008).

GAIA AND HUMANITY: REVERSING DEAD-END HUMAN BEHAVIOR

The essence of Gaian thinking is that the complex living and physical systems of the Earth work together as an integrated whole, with feedback mechanisms within and between systems resulting in stability or recovery from perturbations that are not excessive. Gaia theory expands on this to deduce that life is an *active* participant in this process, by contributing profoundly to the maintenance of dynamic stability of natural systems within ranges necessary for life itself.

Human beings have, within the past 100 years, become the greatest single impacting form of life on Earth. A 2007 analysis, for example, revealed that humans now consume almost 25 percent of the Earth's *total biological productivity*.[28] Consequently the active participation of the life element of Gaia is now largely being determined by the present and near-future actions of the human species. All available evidence today suggests that many human actions are introducing physical and chemical forcings that are too great for the Earth system to counter, thereby leading to ever-greater destabilizations.

The geological record, and especially climatic evidence, indicate that Gaia can be dangerously fast in destabilizing, and maddeningly slow in recovering, This adds an enormous urgency to all nations, religions, and cultures to redefine their most basic social institutions in terms of meeting our collective global responsibility toward a newly stable, biologically rich and diverse Gaia that can support human civilizations for millennia to come.

While the challenges of climate change have been the subject matter of this article, the broad picture of our adverse impact on the Earth system is more extensive: human overpopulation and unsustainable land-use patterns; depletion of freshwater resources; wars (which are terribly destructive of environments, to say nothing of people and cultures); rapidly expanding extinctions, forest destruction, and overall diminishment

of biodiversity; loss of topsoil and decline in soil fertility; air, soil and water pollution; chemical contaminations of almost all living things, including human bodies; and overconsumption of the Earth's resources.

Ironically, this litany suggests that people are knowingly (albeit not willfully) disrupting the Earth system in such a way as to make human life more difficult, and perhaps even untenable in the long run. This does not contradict Gaia theory. Life on Earth can persist in the future in a new equilibrium of physical and living systems, featuring ecosystems and biomes that are different from those of today—without human beings. There is nothing in Gaia theory, after all, that says that life on Earth must include humans.

Humanity now bears the burden of determining whether the Gaia stabilizations necessary for the continued support of the human species can be protected and recovered. This awareness is only now dawning, driven by dramatic and disturbing changes in local and global climate and water systems, changes that are known to be largely the result of human decisions and actions. But humans have *elected to cause* these changes. Why are we being so obtuse?

Because, paradoxically, it seems that we—the ostensible pinnacle of evolutionary development—have invented religious, political, and economic institutions that have caused humans to adopt policies and to undertake actions based on preference, belief, or expediency. The impact on the Earth's living and physical systems, or on future generations, has rarely been considered. As a consequence almost all actions of societies and civilizations over the past couple of centuries—a blip in geological time—have conspired to destroy Gaian life-support systems. Only recently have some societies accepted a measure of this responsibility through emerging environmental protections.

Previously the scale of human impact was small in the great scheme of the Earth's interconnected living and nonliving systems. Now it is suddenly too great. Previously human social, political, cultural, and religious institutions had little effect on Gaia's global mechanisms. Suddenly the effects are dominant.

These same social institutions, however, can provide the tools that must be used to reverse these dangerous directions. This is fortunate, for while the Earth's physical processes possess great inertia, generally moving and changing slowly, humans can elect to change their social structures and adopt new global responsibilities on much shorter time scales. Restoring the Earth is now our responsibility. We need to understand that the path to global re-dependence on the naturally occurring renewable energy sources is not only desirable, it is entirely feasible. Furthermore it has become imperative. Just as we dream of a future without wars over resources, we can share a vision of a future powered by the sun. And what we can dream, we can accomplish.

NOTES

1. Arctic statistics quoted here are from the *Arctic Climate Impact Assessment*, drawn from an overview lecture by Robert Correll. The complete report can be downloaded at http://www.acia.uaf.edu/pages/scientific.html.
2. Inland glaciers are also melting. During the summer 2003 heat wave in Europe about 10 percent of the Alpine glaciers melted. At this pace 75 percent of the glaciers in Switzerland will be gone by 2050. By 2030 Glacier National Park in the United States, will only have memories where the glaciers once were.
3. United Nations General Assembly, 96th Plenary meeting, Resolution 42/186, December 11, 1987, http://www.un.org/documents/ga/res/42/ares42-186.htm.
4. The full text of the UNFCCC United Nations Framework Convention on Climate Change can be found at http://unfccc.int/resource/docs/convkp/conveng.pdf.
5. Kyoto Protocol to the United Nations Framework Convention on Climate Change, United Nations 1998, http://unfccc.int/resource/docs/convkp/kpeng.pdf.
6. An entire conference was devoted to the problem of "Avoiding Dangerous Climate Change" in 2005, producing a book of peer-reviewed papers that substantiate the reasonableness of these numerical targets. The 16.3 MB book can be downloaded from http://www.defra.gov.uk/environment/climatechange/Internat/pdf/avoid-dangercc.pdf.

7. A short and readable summary of these points can be found in *Global Warming: A Target for U.S. Emissions* by the Union of Concerned Scientists, at http://www.ucsusa.org/global_warming/science/emissionstarget.html.
8. For many resources on this subject, see the Pew Center for Global Climate Change, Business and Climate reports, at http://www.pewclimate.org/companies_leading_the_way_belc/business_resource_portal/business_reports_res.cfm.
9. A running tally of cities that have signed onto the agreement is kept on the home page of the US Conference of Mayors, http://www.usmayors.org/uscm/home.asp.
10. http://www.cities-localgovernments.org/uclg/upload/news/newsdocs/World_Mayors_Local_Governments_Climate_Protection_Agreement.pdf.
11. See http://www.architecture2030.org/home.html.
12. US Energy Information Agency. See http://www.architecture2030.org/home.html.
13. Edward Mazria, founder of the Architecture 2030 Challenge, personal communication.
14. See also Pacala and Socolow (2006).
15. The world emits about 7 billion tons of carbon per day. This represents about 25 billion tons of CO_2 equivalent, with the ratio of 3.67 tons of CO_2 for each ton of carbon.
16. The US Nuclear Regulatory Commission is anticipating fast-track license applications by 2009 for 28 new nuclear reactors to be built at 19 sites in the United States.
17. The International Atomic Energy Agency (IAEA) in its October 2007 "high projection" suggests that world nuclear capacity could increase by 84 percent, up to 679GW(e), which would have nuclear on track to accomplish one full wedge by 2050.
18. Coal, uranium 235, petroleum, natural gas, and tar sands. Figures from Steven Heckeroth and Richard Perez.
19. One terawatt hour is one billion kilowatt hours, or 1,000 gigawatt hours. Global consumption of stored energy is about 80,000 terawatt hours/year.
20. The global statistics used in this section are from Martinot (2008).
21. In Germany the renewable energy industries, which produce about 6 percent of Germany's electricity, have created more jobs than the nuclear power industry, which produces about 30 percent of Germany's electricity, for about a 5:1 ratio in job-producing advantages by the renewables.
22. Martinot (2008).
23. This policy concept was co-introduced by the author of this chapter in California Public Utility proceedings in 1995, acting on behalf of the Union of Concerned Scientists, in partnership with Nancy Rader, acting on behalf of the American Wind Energy Association.
24. But the carbon reduction benefits of this modest goal will be completely swamped if China continues to build the 800 coal-fired power plants they are anticipating.
25. See *Tackling Climate Change in the U.S.: Potential Carbon Emissions Reductions from Energy Efficiency and Renewable Energy by 2030*, www.ases.org/climatechange. A summary of the conclusions can be found in Kutscher (2007).
26. An even more ambitious plan for the United States was published in early 2008, showing a purely solar energy path that could provide 69 percent of the country's electricity and 35 percent of the total energy by 2050 (Zweibel et al. 2008). The total cost to accomplish this solar transition for the United States was estimated to be $400 billion over 40 years (the total 40-year cost represents two *years* of the Iraq war costs).
27. Former vice president Al Gore, co-recipient with the IPCC scientists of the 2007 Nobel Peace Prize for their joint work on climate change, and the publicizing of the enormity of the problem, announced in 2008 that a worthy world goal would be to strive for no fossil-fuel use within ten years.
28. *Science News* 172: 235.

REFERENCES

Archer, C. L., and M. Z. Jacobson. 2005. Evaluation of global wind power. *Journal of Geophysical Research* 110: D12110.

Bell, R. E. 2008. The unquiet ice. *Scientific American* 298 (2): 60–67.

David, T. H., G. Lean, and S. Mesur. 2008. Big business says addressing climate change "rates very low on agenda." *The Independent*, January 27.

Emanual, K., R. Sundararajan, and J. Williams. 2008. *Bulletin of the American Meteorological Society* 89 (3): 347–67.

Hanna, E, P. Huybrechts, K. Steffen, J. Cappelen, R. Huff, C. Shuman, T. Irvine-Flyng, and M. Griffiths. 2008. Increased runoff from melt from the Greenland ice sheet: A response to global warming. *Journal of Climate* 21 (2): 331–41.

Hansen, J. 2008, Global warming 20 years later: Tipping points near. Presentation before the National Press

Club and the House Select Committee on Energy Independence and Global Warming, June 23.
Hultman, N. E., J. G. Koomey, and D. M. Kammen. 2007. What history can teach us about the future costs of U.S. nuclear power. *Environmental Science and Technology* 41 (7): 2088–93.
IPCC. 2007. http://www.ipcc.ch/.
Juncosa, B. 2008. Suffocating seas. *Scientific American* 299 (4): 20–22.
Kerr, R. A. 2006. A worrying trend of less ice, higher seas. *Science* 311 (5768): 1698–1701.
Kutscher, C. 2007. Tackling climate change in the U.S. *Solar Today* 21 (2): 26–29.
Lempinen, E. W. 2007. Science and environment: searching for climate change's "tipping points." *Science* 318 (5855): 1396–97.
Mann, M. E., and K. Emanual. 2006. Atlantic hurricane trends linked to climate change. *EOS* 87 (24): 233–44.
Martinot, E. 2008. *Renewables 2007—Global Status Review*. Washington, DC: Worldwatch Institute.
Meinshausen, M. 2006. What does a 2°C target mean for greenhouse gas concentrations? In H. Schellnhuber et al., eds., *Avoiding Dangerous Climate Change*. Cambridge: Cambridge University Press.
Pacala, S., and R. Socolow. 2004. Stabilization wedges: Solving the climate problem for the next 50 years with current technologies. *Science* 305 (5686): 968–72.
Pacala, S., and R. Socolow. 2006. A plan to keep carbon in check. *Scientific American* 297 (9): 50–57.
Perkins, S. 2007. Portrait of a Meltdown. *Science News* 172 (December 29): 387.
Webster, P. J., G. J. Holland, J. A. Curry, and H.-R. Chang. 2005. Changes in tropical cyclone number, duration, and intensity in a warming environment. *Science* 309 (5742): 1844–46.
Zweibel, K., J. Mason, and V. Fthenakis. 2008. A solar grand plan. *Scientific American* 299 (January).

2.6.3 NINE CHALLENGES OF ALTERNATIVE ENERGY

DAVID FRIDLEY

The increasing price of gasoline and home heating oil, along with concerns about the security of our country's energy supply and possible climate change from the burning of fossil fuels, have been major topics of discussion in recent years. Along with these factors, the growing energy needs of developing nations and regions has increased interest in alternative energy sources.

In this selection, David Fridley, a fellow of the Post Carbon Institute with thirty years of experience in the energy field, reviews what he sees as the major issues related to a range of alternative energy paths. Fridley argues that focusing solely on the initial costs of substituting a renewable energy system for the current carbon-based one ignores a number of complex and interrelated issues. These include the viability, scalability, feasibility, and sustainability of various alternative energy approaches, as well as political realities, vested interests, and other relevant conditions. These issues suggest that any significant change of energy policy will only come about if the magnitude of our energy problem is recognized by a majority of the population.

Each of the nine alternative energy challenges Fridley discusses are interdependent to varying degrees. They demonstrate the extent of the problem we face in making what many believe is a necessary shift away from an established and entrenched carbon-based economy. Whether we like it or not, the centuries-long era of cheap fossil fuel energy may soon be over. In Fridley's view, a sustainable, post-carbon future demands that we become more attentive to how we use and conserve many different forms of energy.

Source: From *The Post Carbon Reader* edited by Richard Heinberg and Daniel Lerch. Copyright © 2010 by Post Carbon Institute. Reprinted by permission of Watershed Media.

FOCUS QUESTIONS

1. Why is it not appropriate to compare only the "costs" of various forms of energy? What issues are often overlooked or ignored in attempting to do that?
2. Many argue that we cannot "afford" to shift from our current fossil fuel economy to one that relies more heavily on renewable energy. Why, according to Fridley, is this shift more "affordable" than continuing our current system, and what are the costs of not making it?
3. What ethical issues do you see in the imbalance of per capita energy consumption between the developed and less-developed areas of the world? Is it possible or desirable for underdeveloped regions to start consuming energy at a pace comparable to that of developed regions? What are some different ways that the developing world might gain access to this energy? Which approach would be best? Which would be worst? Please explain your answers with specific reference to Fridley's article.

KEYWORDS

baseload, biomass, bioreactors, capacity factor, conversion efficiency, corn ethanol, deployment, energy bounty, energy return on investment (EROI), feedstocks, footprint, full-cycle water requirement, Hirsch report, hybrid vehicle, input costs, intermittency, megajoules (MJ), net energy, oil shale, photon-absorbing silicon panels, physical accounting, post-carbon world, rare-earth elements, scalability, smart-grid, storage, substitutability, tar sands, thin-film solar, timing

The scramble for alternatives is on. High oil prices, growing concerns over energy security, and the threat of climate change have all stimulated investment in the development of alternatives to conventional oil. "Alternative energy" generally falls into two categories:

- Substitutes for existing petroleum liquids (ethanol, biodiesel, bio-butanol, dimethyl ether, coal-to-liquids, tar sands, oil shale), both from biomass and fossil feedstocks.
- Alternatives for the generation of electric power, including power-storage technologies (wind, solar photovoltaics, solar thermal, tidal, biomass, fuel cells, batteries).

The technology pathways to these alternatives vary widely, from distillation and gasification to bioreactors of algae and high-tech manufacturing of photon-absorbing silicon panels. Many are considered "green" or "clean," although some, such as coal-to-liquids and tar sands, are "dirtier" than the petroleum they are replacing. Others, such as biofuels, have concomitant environmental impacts that offset potential carbon savings.

Unlike conventional fossil fuels, where nature provided energy over millions of years to convert biomass into energy-dense solids, liquids, and gases—requiring only extraction and transportation technology for us to mobilize them—alternative energy depends heavily on specially engineered equipment and infrastructure for capture or conversion, essentially making it a high-tech manufacturing process. However, the full supply chain for alternative energy, from raw materials to manufacturing, is still very dependent on fossil-fuel energy for mining, transport, and materials production. Alternative energy faces the challenge of how to supplant a fossil-fuel-based supply chain with one driven by alternative energy forms themselves in order to break their reliance on a fossil-fuel foundation.

The public discussion about alternative energy is often reduced to an assessment of its monetary costs versus those of traditional fossil fuels, often in comparison to their carbon footprints. This kind of reductionism to a simple monetary metric obscures the complex issues surrounding the potential viability, scalability, feasibility, and suitability of pursuing specific alternative technology paths. Although money is necessary to develop alternative energy, money is simply a token for mobilizing a range of resources used to produce energy. At the level of physical requirements, assessing the potential for alternative energy development becomes much more

complex since it involves issues of end-use energy requirements, resource-use trade-offs (including water and land), and material scarcity.

Similarly, it is often assumed that alternative energy will seamlessly substitute for the oil, gas, or coal it is designed to supplant—but this is rarely the case. Integration of alternatives into our current energy system will require enormous investment in both new equipment and new infrastructure—along with the resource consumption required for their manufacture—at a time when capital to make such investments has become harder to secure. This raises the question of the suitability of moving toward an alternative energy future with an assumption that the structure of our current large-scale, centralized energy system should be maintained. Since alternative energy resources vary greatly by location, it may be necessary to consider different forms of energy for different localities.

It is not possible to single out one metric by which to assess the promise of a particular alternative energy form. The issue is complex and multifaceted, and its discussion is complicated by political biases, ignorance of basic science, and a lack of appreciation of the magnitude of the problem. Many factors come into play, of which nine are discussed here.

1. SCALABILITY AND TIMING

For the promise of an alternative energy source to be achieved, it must be supplied in the time frame needed, in the volume needed, and at a reasonable cost. Many alternatives have been successfully demonstrated at the small scale (algae-based diesel, cellulosic ethanol, biobutanol, thin-film solar) but demonstration scale does not provide an indication of the potential for large-scale production. Similarly, because alternative energy relies on engineering and construction of equipment and manufacturing processes for its production, output grows in a stepwise function only as new capacity comes online, which in turn is reliant on timely procurement of the input energy and other required input materials. This difference between "production" of alternative energy and "extraction" of fossil fuels can result in marked constraints on the ability to increase the production of an alternative energy source as it is needed.

For example, the tar sands of Canada (although often excluded as an "alternative" energy, tar sands are subject to the same constraints because the production of oil from the tar-sands deposits is essentially a mining and manufacturing operation) have already achieved a fully commercial scale of production, and because of the immense reserves indicated in Alberta, tar sands are looked to be a backstop to declining conventional crude oil production. In 2008, production of oil from the tar sands reached 1.2 million barrels per day (bpd), less than 2 percent of global production of conventional crude oil. By 2020, the Canadian Association of Petroleum Producers projects that production will increase by 2.1 million bpd to a total of 3.3 million bpd.[1] But the International Energy Agency (IEA) estimates that the global decline rate from conventional-oil fields is 6.4 percent, or about 4.8 million bpd per year.[2] Thus by 2020, the new oil coming from tar-sands production will not even make up half of what is being lost from ongoing depletion of existing conventional-oil fields. Even with a "crash" production program, it is estimated that tar-sands production in 2020 could not exceed 4.0 million bpd, an increase still less than the annual rate of conventional crude oil depletion.[3]

Scale also matters in comparing projected production of an alternative energy form against expected demand growth. In 2007, the U.S. Energy Policy Act established a target for the production of ethanol in 2022 at 36 billion gallons, of which 15 billion gallons were to be sourced from corn and the remainder from cellulosic sources. In terms of gasoline equivalency, this target is equal to 890,000 bpd of additional supply. In 2008, however, the U.S. Department of Energy, in its *Annual Energy Outlook*, forecast demand for gasoline would grow by 930,000 bpd by 2022,[4] more than offsetting projected supply growth from ethanol and leaving gross oil dependency unchanged.

This lack of the kind of scalability needed given the magnitude and time frame of conventional-oil depletion and in the face of continued demand growth is found as well in other biofuels, coal-to-liquids, and alternative liquids for transportation. Also of concern is the difficulty of scaling up alternative energy quickly enough to meet greenhouse gas emissions targets.

2. COMMERCIALIZATION

Closely related to the issue of scalability and timing is commercialization, or the question of how far away a proposed alternative energy source stands from being fully commercialized. Often, newspaper reports of a scientific laboratory breakthrough are accompanied by suggestions that such a breakthrough represents a possible "solution" to our energy challenges. In reality, the average time frame between laboratory demonstration of feasibility and full large-scale commercialization is twenty to twenty-five years. Processes need to be perfected and optimized, patents developed, demonstration tests performed, pilot plants built and evaluated, environmental impacts assessed, and engineering, design, siting, financing, economic, and other studies undertaken. In other words, technologies that are proved feasible on the benchtop today will likely have little impact until the 2030s. This reality is reflected in the key message of the now-famous Hirsch Report, which noted that to properly mitigate the economic impacts of peak oil [2005], we would have needed to start fundamentally redesigning our national energy infrastructure *twenty years* in advance of the peak.[5]

3. SUBSTITUTABILITY

Ideally, an alternative energy form would integrate directly into the current energy system as a "drop-in" substitute for an existing form without requiring further infrastructure changes. This is rarely the case, and the lack of substitutability is particularly pronounced in the case of the electrification of transportation, such as with electric vehicles. Although it is possible to generate the electricity needed for electrified transportation from wind or solar power, the prerequisites to achieving this are extensive. Electric-car development would require extensive infrastructure changes, including:

Retooling of factories to produce the vehicles

- Development of a large-scale battery industry
- Development of recharging facilities
- Deployment of instruments for the maintenance and repair of such vehicles
- A spare-parts industry
- "Smart-grid" monitoring and control software and equipment
- Even more generation and transmission facilities to supply the additional electricity demand

The development of wind and solar-power electricity also requires additional infrastructure; wind and solar electricity must be generated where the best resources exist, which is often far from population centers. Thus, extensive investment in transmission infrastructure to bring it to consumption centers is required. Today, ethanol can be blended with gasoline and used directly, but its propensity to absorb water and its high oxygen content make it unsuitable for transport in existing pipeline systems,[6] and an alternative pipeline system to enable its widespread use would be materially and financially intensive. While alternative energy forms may provide the same energy services as another form, they rarely substitute directly, and these additional material costs need to be considered.

4. MATERIAL INPUT REQUIREMENTS

Unlike what is generally assumed, the input to an alternative energy process is not money per se: It is resources and energy, and the type and volume of the resources and energy needed may in turn limit the scalability and affect the cost and feasibility of an alternative. This is particularly notable in processes that rely on advanced technologies manufactured with rare-earth elements. Fuel cells, for example, require platinum, palladium, and rare-earth elements. Solar-photovoltaic technology requires gallium, and in some forms, indium. Advanced batteries rely on lithium. Even technology designed to save energy, such as light-emitting diode (LED) or organic LED (OLED) lighting, requires rare earths, indium, and gallium. Expressing the costs of alternative energy only in monetary terms obscures potential limits arising from the requirements for resources and energy inputs.

Because alternative energy today constitutes only a small fraction of total energy production, the volume of resources and energy demanded for its production has so far been easily accommodated. This will not necessarily be the case with

large-scale expansion. For example, thin-film solar has been promoted as a much lower-cost, more flexible, and more widely applicable solar-conversion technology compared to traditional silicon panels. Thin-film solar currently uses indium because of its versatile properties, but indium is also widely used as a component of flat-screen monitors. Reserves of indium are limited, and a 2007 study found that at current rates of consumption, known reserves of indium would last just thirteen years.[7]

Can greatly increased demand for these resources be accommodated? As shown in Table 1, successful deployment to 2030 of a range of new energy technologies (and some non-energy advanced technologies) would substantially raise demand for a range of metals beyond the level of world production today. In the case of gallium, demand from emerging technologies would be expected to reach six times today's total global production by 2030; for indium, more than three times today's production—compared to just fractional increases in the demand for ruthenium and selenium.

Although alternative metals and materials exist for certain technologies (albeit often with performance trade-offs), embarking on a particular technology deployment path without consideration of long-term availability of material inputs can substantially raise risks. These risks are not limited to physical availability and price; they include potential supply disruptions as

TABLE 1 Global Demand on Raw Materials from Emerging Technologies

RAW MATERIAL	FRACTION OF TODAY'S TOTAL WORLD PRODUCTION		EMERGING TECHNOLOGIES (SELECTED)
	2006	2030	
Gallium	0.28	6.09	Thin-layer photovoltaics, integrated circuits, white LEDs
Neodymium	0.55	3.82	Permanent magnets, laser technology
Indium	0.40	3.29	Displays, thin-layer photovoltaics
Germanium	0.31	2.44	Fiber-optic cable, infrared optical technologies
Scandium	Low	2.28	Solid oxide fuel cells, aluminum alloying element
Platinum	Low	1.56	Fuel cells, catalysts
Tantalum	0.39	1.01	Microcapacitors, medical technology
Silver	0.26	0.78	Radio-frequency ID tags, lead-free soft solder
Tin	0.62	0.77	Lead-free soft solder, transparent electrodes
Cobalt	0.19	0.40	Lithium-ion batteries, synthetic fuels
Palladium	0.10	0.34	Catalysts, seawater desalination
Titanium	0.08	0.29	Seawater desalination, implants
Copper	0.09	0.24	Efficient electric motors, radio-frequency ID tags
Selenium	Low	0.11	Thin-layer photovoltaics, alloying element
Niobium	0.01	0.03	Microcapacitors, ferroalloys
Ruthenium	0.00	0.03	Dye-sensitized solar cells, Ti-alloying element
Yttrium	Low	0.01	Superconduction, laser technology
Antimony	Low	Low	Antimony-tin-oxides, microcapacitors
Chromium	Low	Low	Seawater desalination, marine technologies

Source: Gerhard Angerer et al., "Raw Materials for Emerging Technologies," (Karlsruhe: Fraunhofer Institute for Systems and Innovation Research ISI; Berlin: Institute for Futures Studies and Technology Assessment IZT, February 2, 2009).

a consequence of the uneven geographical distribution of production and reserves. Currently, China is the dominant world source (over 95 percent) of the rare-earth element neodymium, a key input in the production of permanent magnets used in hybrid-vehicle motors and windmill turbines. In 2009, the Chinese government announced restrictions on the export of rare earths, ostensibly to encourage investment within China of industries using the metals. Whether for the rare earths themselves or for final products made from them, import dependency in the face of such a high concentration of production would do little to alleviate energy security concerns now seen in terms of import dependency on the Middle East for oil.

Alternative energy production is reliant not only on a range of resource inputs, but also on fossil fuels for the mining of raw materials, transport, manufacturing, construction, maintenance, and decommissioning. Currently, no alternative energy exists without fossil-fuel inputs, and no alternative energy process can reproduce itself—that is, manufacture the equipment needed for its own production—without the use of fossil fuels. In this regard, alternative energy serves as a supplement to the fossil-fuel base, and its input requirements may constrain its development in cases of either material or energy scarcity.

5. INTERMITTENCY

Modern societies expect that electrons will flow when a switch is flipped, that gas will flow when a knob is turned, and that liquids will flow when the pump handle is squeezed. This system of continuous supply is possible because of our exploitation of large stores of fossil fuels, which are the result of millions of years of intermittent sunlight concentrated into a continuously extractable source of energy. Alternative energies such as solar and wind power, in contrast, produce only intermittently as the wind blows or the sun shines, and even biomass-based fuels depend on seasonal harvests of crops. Integration of these energy forms into our current system creates challenges of balancing availability and demand, and it remains doubtful that these intermittent energy forms can provide a majority of our future energy needs in the same way that we expect energy to be available today.

One indication of intermittency challenges in electric power generation is the capacity factor, or the average percentage of time in a year that a power plant is producing at full rated capacity. As shown in Table 2, photovoltaic systems produce at full capacity only 12 to 19 percent of the time over the course of a year, compared to an average of 30 percent for wind systems. In contrast, a coal-thermal plant will typically run at full capacity 70 to 90 percent of the time, while nuclear power operates at over a 90 percent capacity factor in the United States.

Our current electricity system is dominated by large baseload coal-and nuclear-power generation. The integration of intermittent energy forms such as solar and wind is increasingly seen as a matter of expanding transmission capacity and grid interconnections to extend the area over which these variations are felt, as well as implementing more complex operations controls. This approach in effect relies on strengthening and expanding the large centralized energy production and

TABLE 2 Common Capacity Factors for Power Generation

GENERATION TYPE	CAPACITY FACTOR
Photovoltaics	12–19%
Thermal solar	~15%
Thermal solar with storage	70–75%
Wind	20–40%
Hydropower	30–80%
Geo thermal	70–90%
Nuclear power	60–100%
Natural gas combined cycle (non-peaker)	~60%
Coal thermal	70–90%

Sources: Renewable Energy Research Laboratory, University of Massachusetts at Amherst, *Wind Power: Capacity Factor, Intermittency*; National Renewable Energy Laboratory, *Assessment of Parabolic Trough and Power Tower Solar Technology Cost and Performance Forecast*, NREL/SR-550-34440 (Golden, CO: NREL, 2003).

distribution model that has characterized the fossil-fuel era, but may not necessarily be suitable for a future of renewable energy generation.

The key to evening out the impact of intermittency is storage; that is, the development of technologies and approaches that can store energy generated during periods of good wind and sun for use at other times. Many approaches have been proposed and tested, including compressed-air storage, batteries, and the use of molten salts in solar-thermal plants.

The major drawbacks of all these approaches include the losses involved in energy storage and release, and the limited energy density that these storage technologies can achieve.

6. ENERGY DENSITY

Energy density refers to the amount of energy that is contained in a unit of an energy form. It can be expressed in the amount of energy per unit of mass (weight) or in the amount of energy per unit of volume. In everyday life, it is common to consider energy density when considering food choices. Food labeling in the United States requires that both numbers needed for calculating energy density be provided: the number of food calories per serving and the weight or volume of the serving (expressed in grams or liters, respectively). Potatoes, for example, have an energy density of 200 food calories per 100 grams, or, expressed in units common in energy discussions, 8.4 megajoules[8] (MJ) per kilogram (about 2.2 pounds). Cheese is more energy dense than potatoes, containing about 13 MJ per kilogram.

Energy density has also influenced our choice of fuels. Aside from alleviating a growing wood shortage, the conversion to the use of coal in the seventeenth and eighteenth centuries was welcomed because coal provided twice as much energy as wood for the same weight of material. Similarly, the shift from coal- to petroleum-powered ships in the early-twentieth century was driven by the fact that petroleum possesses nearly twice the energy density of coal, allowing ships to go farther without having to stop for refueling. Even when used in a motor vehicle's inefficient internal combustion engine, a kilogram of highly energy-dense gasoline—about 6 cups—allows us to move 3,000 pounds of metal about 11 miles.

The consequence of low energy density is that larger amounts of material or resources are needed to provide the same amount of energy as a denser material or fuel. Many alternative energies and storage technologies are characterized by low energy densities, and their deployment will result in higher levels of resource consumption. As shown in Figure 1, the main alternatives under development to supplant gasoline use in cars are dramatically lower in energy density than gasoline itself. Lithium-ion batteries—the focus of current research for electric vehicles—contain only 0.5 MJ per kilogram of battery compared to 46 MJ per kilogram for gasoline. Advances in battery technology are being announced regularly, but they all come up against the theoretical limit of battery density of only 3 MJ per kilogram. Low energy density will present a significant challenge

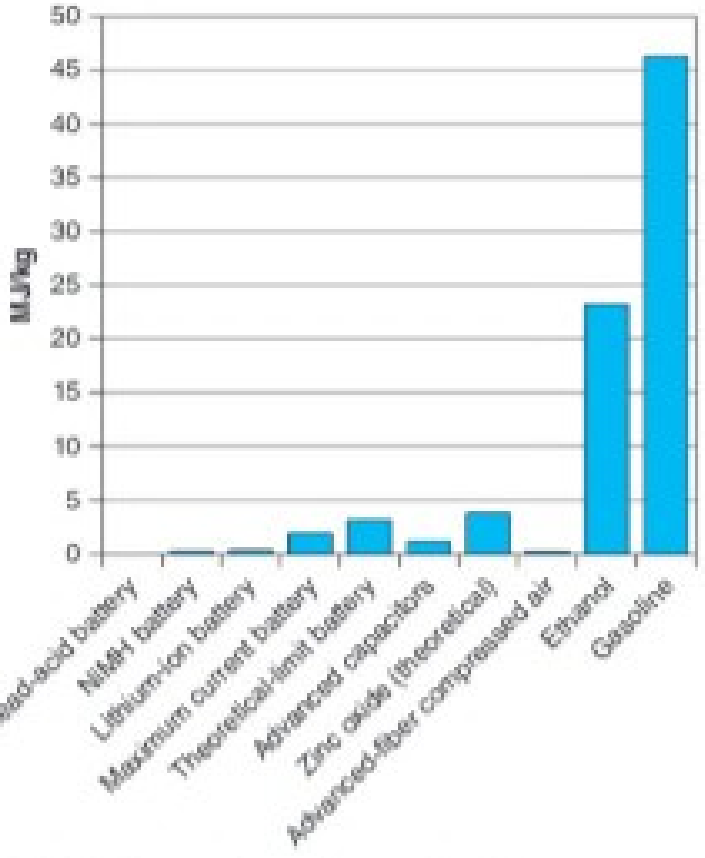

FIGURE 1 Comparison of energy densities.

Source: Kurt Zenz House and Alex Johnson, "The Limits of Energy Storage Technology," *Bulletin of the Atomic Scientists* Web edition, January 20, 2009, http://www.thebulletin.org/web-edition/columnists/kurt-zenz-house/the-limits-of-energy-storage-technology.

to the electrification of the car fleet and will raise challenges of adequate material supply: Today, the advanced Tesla Roadster has a lithium-ion battery pack weighing 900 pounds, which delivers just 190 MJ of energy. In contrast, a 10-gallon tank of gasoline weighs 62 pounds and delivers 1,200 MJ of energy. To provide the equivalent energy to a typical gasoline car, an electric-car battery pack would need to consume resources weighing 5,700 pounds, nearly the weight of the last Hummer model.

The more dense an energy form is, the less land is needed for its deployment. Because many alternative energies are far less energy dense than fossil fuels, large-scale deployment will incur considerable land costs. For example, a single 1,000-megawatt coal-fired power plant requires 1 to 4 square kilometers (km^2) of land, not counting the land required to mine and transport the coal. In contrast, 20–50 km^2, or the size of a small city, would be required to generate the equivalent amount of energy from a photovoltaic array or from a solar-thermal system. For wind, 50–150 km^2 would be needed; for biomass, 4,000–6,000 km^2 of land would be needed. The sprawling city of Los Angeles, in comparison, covers 1,200 km^2. The land-use issue is thus a problem not only of biofuels production; siting of alternative energy projects will likely be a constant challenge because of the inherent high land footprint.

7. WATER

Water ranks with energy as a potential source of conflict among peoples and nations, but a number of alternative energy sources, primarily biomass-based energy, are large water consumers critically dependent on a dependable water supply. As seen in Figure 2, the "full-cycle" water requirement (including water for growing and processing biofuels) for key ethanol and biodiesel feedstocks is in some cases hundreds or even many thousand times higher than for the refining of gasoline. In well-watered regions with regular and adequate rainfall, much of this water can be provided through rain; in a region such as California, where no rain falls during the summer growing season because of its Mediterranean climate, irrigation is an absolute necessity for growing commercial biomass feedstocks. However, all of California's water resources have already been

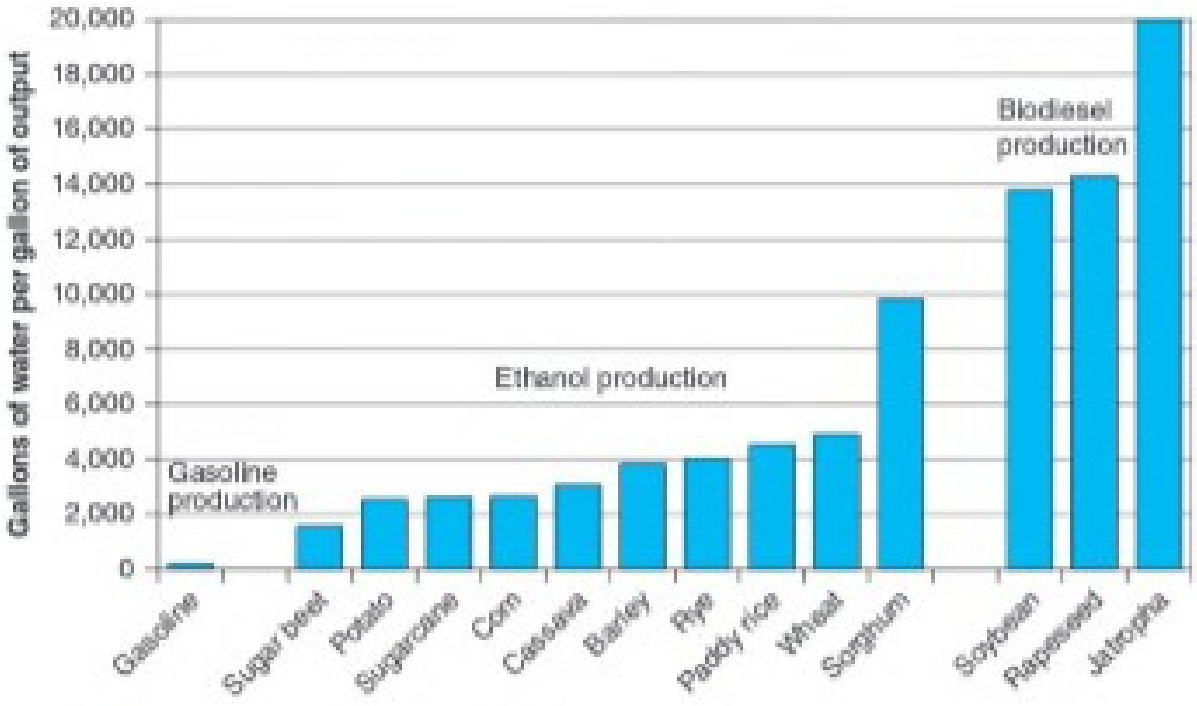

FIGURE 2 Full-cycle water requirements for biofuel production.

Source: Winnie Gerbens-Leenes et al., "The Water Footprint of Bioenergy," *Proceedings of the National Academy of Sciences* 106, no. 25 (June 23, 2009): 10219–10223.

allocated, so existing uses for other crops would have to be reallocated to support biomass farming—raising the issue of "food versus fuel" from yet a different angle. The water problems, however, promise only to intensify with global warming as California's winter snowpack fades and runoff to support summer agriculture declines.

Considering just the processing stage, biomass and unconventional fossil-fuel energy also often require much greater water usage than the 2.5 gallons of water required per gallon of gasoline produced. Coal-to-liquids production consumes 8 to 11 gallons of water per gallon of output, corn ethanol requires 4 to 6 gallons, and cellulosic ethanol needs 11 gallons. In the United States, Montana has looked into becoming a leader in coal-to-liquids production, yet Montana's dry climate suggests that water could be a limiting factor.

8. THE LAW OF RECEDING HORIZONS

An often-cited metric of the viability of alternatives is the expected breakeven cost of the alternative with oil, or the price that crude oil would have to be to make the alternative cost competitive. Underlying this calculation, however, is an assumption that the input costs to alternative energy production would remain static as oil prices rise, thereby providing the economic incentive to development. This assumption, however, has not always proved to be the case, particularly for those alternatives for which energy itself is a major input. Because of price linkages in the energy (and now energy and biomass) markets, rising oil prices tend to push up the price of natural gas as well as coal; for processes that are heavily dependent on these fuels, higher oil prices also bring higher production costs.

A good example of this phenomenon is the assessment of the economics of production from oil shale (kerogen-rich marlstone), found in vast quantities in Colorado, Utah, and Wyoming. In the early 1970s, shale oil was expected to flood the market if the price of crude oil were to rise above $2 per barrel. "When world oil prices had shot up to $35 per barrel by 1979, oil-shale production still required federal government assistance, and when oil prices fell in the mid-1980s, development and production were abandoned. Fast-forward to 2008 when oil prices moved above $100 per barrel—oil shale was then expected to be economic at $80 to $90 per barrel, and the U.S. government again provided incentives to explore production in the area. This ratcheting up of oil-shale economics with the price of oil reflects in part the high energy-input requirement to the production process.

Similarly, the corn ethanol industry has recently been subject to the same dynamic step-up in costs as the price of oil has risen. Two major input costs to the industry are the processing fuel (usually natural gas) and the corn feedstock itself. Rising oil prices after 2004 pulled natural gas prices up as well, increasing the processing energy costs for ethanol. At the same time, higher fuel prices made cultivating corn more expensive because of higher costs for operating farm machinery; this, together with the additional demand for corn created by the growing ethanol industry, helped push corn prices up even further. So, although the record-high oil prices of 2008 increased demand for ethanol, some ethanol producers were operating with minimal or no profit because they had to pay more for both their processing fuel and their corn feedstock.

Ultimately, the "law of receding horizons" is a phenomenon reflective of the general orientation toward financial and economic accounting to gauge project viability and prospects.[9] Physical accounting—that is, analyzing the material and energy inputs to a process—would help in better understanding the degree to which an alternative energy production process is vulnerable to the rise in energy costs.

9. ENERGY RETURN ON INVESTMENT

The complexity of our economy and society is a function of the amount of net energy we have available. "Net energy" is, simply, the amount of energy remaining after we consume energy to produce energy. Consuming energy to produce

energy is unavoidable, but only that which is not consumed to produce energy is available to sustain our industrial, transport, residential, commercial, agricultural, and military activities. The ratio of the amount of energy we put into energy production and the amount of energy we produce is called "energy return on investment" (EROI).[10]

This concept differs from "conversion efficiency," which compares the amount of energy provided as a feedstock to a conversion process (such as an electric power plant or petroleum refinery) with the amount remaining after conversion. Physics dictates that this figure is always less than 100 percent. In contrast, EROI can be very high (e.g., 100:1, or 100 units of energy produced for every 1 unit used to produce it—an "energy source") or low (0.8:1, or only 0.8 unit of energy produced for every 1 unit used in production—an "energy sink"). Society requires energy sources, not energy sinks, and the magnitude of

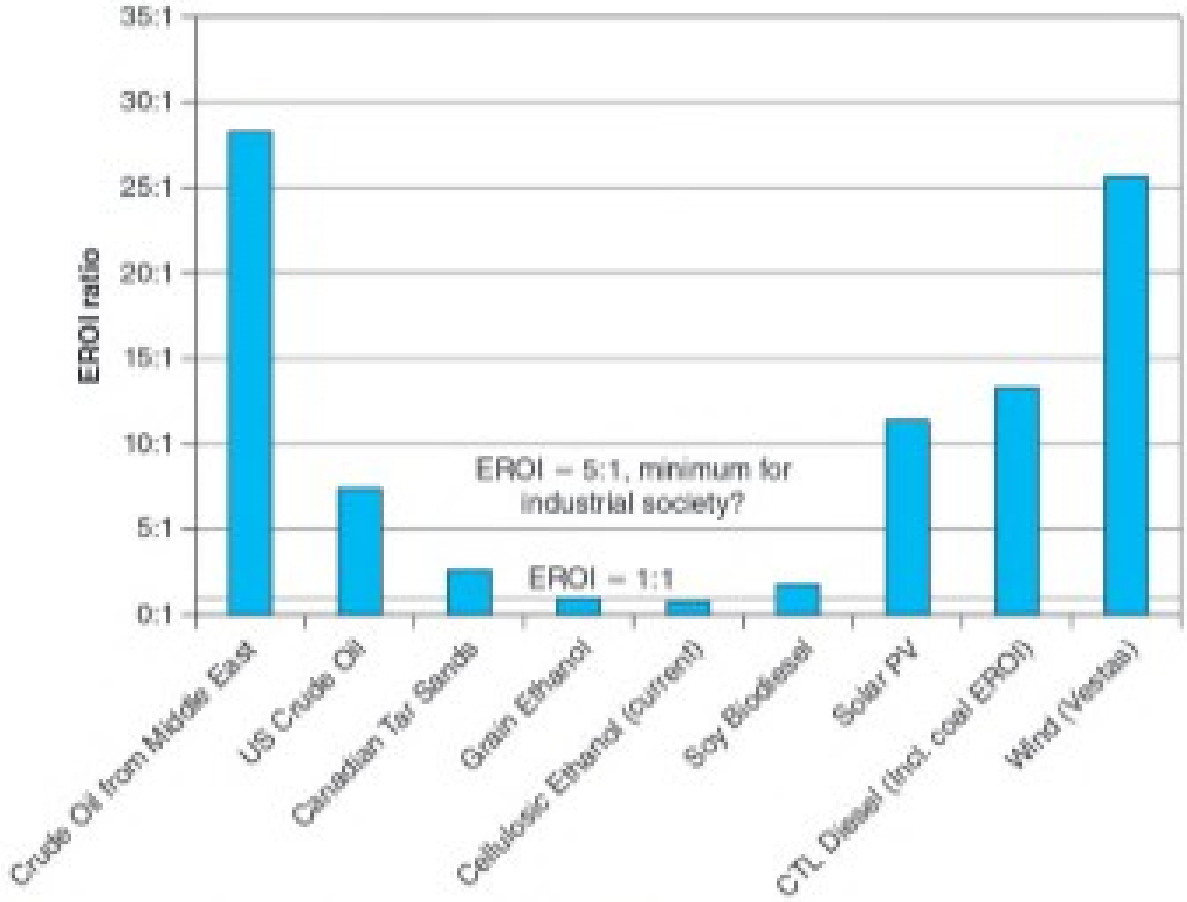

FIGURE 3 Estimated EROI of selected alternatives.

Note: EROI measurements are not standardized; the shading indicates ranges from various studies. PV = photovoltaic; CTL = coal-to-liquids; Vestas = Vestas Wind Systems.

Sources: P. J. Meier and G. L. Kulcinski, *Life-Cycle Energy Requirements and Greenhouse Gas Emissions for Building-Integrated Photovoltaics*, Fusion Technology Institute (2002); National Renewable Energy Laboratory, *What Is the Energy Payback for PV?* DOE/ GO-102004-1847 (2004); Vasilis Fthenakis and Erik Alsema, "Photovoltaics Energy Payback Times, Greenhouse Gas Emissions and External Costs," *Progress in Photovoltaics* 14, no. 3 (May 2006), 275–280; Luc Gagnon, "Civilisation and Energy Payback," *Energy Policy* 36, no. 9 (September 2008), 3317–3322; Cutler Cleveland, "Net Energy from the Extraction of Oil and Gas in the United States, 1954–1997," *Energy* 30 (2005), 769–782; Charles A. S. Hall et al., "Peak Oil, EROI, Investments and the Economy in an Uncertain Future," in *Biofuels, Solar and Wind as Renewable Energy Systems*, David Pimentel, ed. (Springer Science, 2008); Vestas Wind Systems, *Life Cycle Assessment of Offshore and Onshore Sited Wind Power Plants Based on Vestas V90-3.0 MW Turbines* (Denmark, 2006); Alexander E. Farrell et al., "Ethanol Can Contribute to Energy and Environmental Goals," *Science* 311 (January 27, 2006).

EROI for an energy source is a key indicator of its contribution to maintenance of social and economic complexity.

Net-energy availability has varied tremendously over time and in different societies. In the last advanced societies that relied only on solar power (sun, water power, biomass, and the animals that depended on biomass), in the seventeenth and early-eighteenth centuries, the amount of net energy available was low and dependent largely on the food surpluses provided by farmers. At that time, only 10 to 15 percent of the population was *not* involved in energy production. As extraction of coal, oil, and natural gas increased in the nineteenth and twentieth centuries, society was increasingly able to substitute the energy from fossil fuels for manual or animal labor, thereby freeing an even larger proportion of society from direct involvement in energy production. In 1870, 70 percent of the U.S. population were farmers; today the figure is less than 2 percent, and every aspect of agricultural production now relies heavily on petroleum or natural gas. The same is true in other energy sectors: Currently, less than 0.5 percent of the U.S. labor force (about 710,000 people) is directly involved in coal mining, oil and gas extraction, petroleum refining, pipeline transport, and power generation, transmission, and distribution.

The challenge of a transition to alternative energy, then, is whether such energy surpluses can be sustained, and thus whether the type of social and economic specialization we enjoy today can be maintained. Indeed, one study estimates that the minimum EROI for the maintenance of industrial society is 5:1, suggesting that no more than 20 percent of social and economic resources can be dedicated to the production of energy without undermining the structure of industrial society.[11]

In general, most alternative energy sources have low EROI values (see Figure 3). Because of their high energy-input requirements, biofuels produce very little or no energy surplus.[12] Similarly, tar sands provide less than 3 units of energy for each unit consumed. In contrast, wind energy shows a high return on energy investment, but it is subject to the problems of intermittency and siting issues.

A high EROI is not sufficient to ensure that the structure of modern society and economies can be maintained, but it is a prerequisite. Unfortunately, EROI is not well understood or routinely used in energy analyses by government or industry, despite the insights it can provide. Because of the enormous investment in resources and energy that any alternative energy pathway will require, it is important that we look beyond simple financial payback, particularly in a future of rising energy prices, declining fossil-fuel resources, and increasing danger of climate catastrophe.

HOW WILL SOCIETY EVOLVE IN A POST-CARBON WORLD?

Alternative energy forms are crucial for a global transition away from fossil fuels, despite the myriad challenges of their development, scaling, and integration. In face of the peaking of global oil production—to be followed by peaks in natural gas and coal extraction—and of the need to reverse trajectory in carbon emissions, alternative energy sources will need to form the backbone of a future energy system.

That system, however, will not be a facsimile of the system we have today based on continuous uninterrupted supply growing to meet whatever demand is placed on it. As we move away from the energy bounty provided by fossil fuels, we will become increasingly reliant on tapping the current flow of energy from the sun (wind, solar) and on new energy manufacturing processes that will require ever larger consumption of resources (biofuels, other manufactured liquids, batteries). What kind of society we can build on this foundation is unclear, but it will most likely require us to pay more attention to controls on energy demand to accommodate the limitations of our future energy supply. Moreover, the modern focus on centralized production and distribution may be hard to maintain, since local conditions will become increasingly important in determining the feasibility of alternative energy

NOTES

1. Canadian Association of Petroleum Producers, *Crude Oil: Forecast, Markets & Pipeline Expansions*, June 2009, http://www.capp.ca/.
2. International Energy Association, *World Energy Outlook* 2008, http://www.worldenergyoutlook.org/.
3. Bengt Soderberg et al., "A Crash Program Scenario for the Canadian Oil Sands Industry," *Energy Policy* 35, no. 3 (March 2007), 1931–1947.
4. U.S. Energy Information Administration, *Annual Energy Outlook 2008 (First Release)*, DOE/EIA-0383 (2008), January 2008. Note that subsequent revisions of *Annual Energy Outlook 2008* changed the cited figures for gasoline demand.
5. Robert L. Hirsch, Roger Bezdek, and Robert Wendling, *Peaking of World Oil Production: Impacts, Mitigation, & Risk Management*, U.S. Department of Energy report, February 2005, http://www.netl.doe.gov/publications/others/pdf/oil_peaking_netl.pdf.
6. John Whims, "Pipeline Considerations for Ethanol," Agricultural Marketing Resource Center, http://www.agmrc.org/media/cms/ksupipelineethl_8BA5CDF1FD179.pdf.
7. David Cohen, "Earth's Natural Wealth: An Audit," *New Scientist* 23 (May 2007), 34–41.
8. A megajoule equals 239 food calories; a typical adult male requires 10 MJ of food energy per day.
9. As coined by user HelsSoFly, comment on "Drumbeat," The Oil Drum, March 7, 2007, http://www.theoildrum.com/node/2344.
10. For a more in-depth discussion of energy return on investment, see Richard Heinberg, *Searching for a Miracle: "Net Energy" Limits & the Fate of Industrial Society* (San Francisco: Post Carbon Institute/International Forum on Globalization, 2009). Note that EROI is sometimes also referred to as "energy returned on energy invested" (EROEI).
11. Charles A. S. Hall, Robert Powers, and William Schoenberg, "Peak Oil, EROI, Investments and the Economy in an Uncertain Future," in *Biofuels, Solar and Wind as Renewable Energy Systems: Benefits and Risks*, David Pimentel, ed. (New York: Springer, 2008), 109–132.
12. The often-cited 8:1 return on Brazilian ethanol and the high return estimated for cellulosic ethanol are not energy calculations; in these studies, the energy provided from biomass combustion is ignored. See, for example, Suani Teixeira Coelho et al., "Brazilian Sugarcane Ethanol: Lessons Learned," *Energy for Sustainable Development* 10, no. 2 (June 2006), 26–39.

2.6.4 STABILIZING CLIMATE: AN ENERGY EFFICIENCY REVOLUTION

LESTER BROWN

We end this book with a selection from the work of Lester R. Brown, noted author and environmentalist and founder of both the Worldwatch Institute and the Earth Policy Institute. Dr. Brown is the author of more than fifty books on global environmental issues and is the recipient of numerous prizes and honors, including twenty-six honorary degrees and a MacArthur "genius" award.

This chapter, from his 2009 book *Plan B 4.0: Mobilizing to Save Civilization*, provides an overview of how new technologies that promote energy savings can help to reduce dependence on fossil fuels, stabilize the level of CO_2 in the atmosphere, and save money. Brown is a "techno-optimist" who believes that a rapid transition to energy-efficient technologies and renewable energy sources is both feasible and

Source: From *Plan B 4.0: Mobilizing to Save Civilization* by Lester R. Brown, Chapter No. 4, 79–108. Copyright © 2009 Earth Policy Institute. Reprinted by permission of W.W. Norton & Company, Inc. Footnotes have been omitted. Interested readers should consult the original.

necessary, but only if we can mobilize the political will to enact forward-looking energy policies. Drawing on examples of such policy initiatives in several countries, Brown argues that we can meet the goal of stabilizing CO_2 in the atmosphere to 400 parts per million (ppm) by the year 2020.

The book from which selection is drawn is available as a free PDF on the Earth Policy Institute Website http://www.earth-policy.org/books/pb4, which also contains Plan B updates, data sets, a slideshow, and information about the documentary film based on this work, *Journey to Earth*.

FOCUS QUESTIONS

1. The movement to replace older incandescent lighting with new energy efficient bulbs is already taking place. What are some indicators that you have observed of this switch to more energy efficient lighting?
2. What are some of the measures that can be undertaken to reduce energy use from household appliances? How many of these measures have been adopted in your own home?
3. What are some of the direct and indirect ways that we can reduce our dependence on energy from petroleum in the transportation sector? What is the role of public policy in helping to make this happen?
4. What are some examples of how recycling materials can help save energy and reduce landfill? Are there any economic incentives that encourage people to recycle? Discuss.
5. Plan B calls for keeping the atmospheric CO_2 levels from exceeding 400 ppm by the year 2020 by offsetting the predicted growth in energy use with a rapid shift from fossil fuels to renewable energy, and by improving energy efficiency. The author cites several reasons why this goal is both possible and economically prudent, but skepticism remains about how feasible this goal really is. Discuss the reasons for both techno-optimism and pessimism related to the goal of stabilizing climate change through changing the energy economy.

KEYWORDS

carbon emissions, CFLs, Clinton Climate Initiative, disassembly, high-speed trains, LEED certification, plug-in hybrid cars, recycling, renewable energy, smart grids, zero-carbon buildings

The world is in the early stages of two energy revolutions. The first is a shift to new energy-efficient technologies across the board. The larger energy savings potentials include shifting from century-old technologies such as incandescent light bulbs and internal combustion engines to far more efficient technologies. Incandescents are being replaced by compact fluorescent bulbs that use one fourth as much electricity. This in turn will be cut in half by the light-emitting diodes (LEDs) coming on the market. And the most advanced plug-in hybrid car prototypes use only one fifth as much gasoline per mile as the average U.S. car on the road today.

The second energy revolution—the shift from an economy powered by oil, coal, and natural gas to one powered by wind, solar, and geothermal energy—is under way and moving fast. In Europe, new electrical generating capacity from wind, solar, and other renewables now exceeds that from fossil fuels by a wide margin. In the United States, new wind-generating capacity of 8,400 megawatts in 2008 dwarfed the 1,400 megawatts from coal. Nuclear power is fading, too. Worldwide, nuclear power generation actually declined in 2008 while wind electric generating capacity increased by 27,000 megawatts, enough to supply

8 million American homes. The world is changing fast.

This chapter begins with a brief description of Plan B's goal of cutting net carbon emissions and then describes in detail the components of the first revolution—the push to raise energy efficiency worldwide.

Implementing Plan B entails cutting net carbon dioxide (CO_2) emissions 80 percent by 2020. This would keep atmospheric CO_2 levels from exceeding 400 parts per million (ppm), up only modestly from 386 ppm in 2008.

This sets the stage for reducing CO_2 concentrations to the 350 ppm that James Hansen and other climate scientists think is needed to avoid runaway climate change. It will also help keep future temperature rise to a minimum. Such a basic economic restructuring in time to avoid catastrophic climate disruption will be challenging, but how can we face the next generation if we do not try?

This restructuring of the world energy economy is being driven by some traditional concerns and some newer ones. Among the former are mounting concerns over climate change, a growing sense of oil insecurity, the rising level and volatility of fossil fuel prices, and financial outlays for importing oil.

The recent global economic downturn and the record number of young people entering job markets in developing countries has also made labor intensity a goal of energy policymaking. Improving energy efficiency and developing renewable sources of energy are both much more labor-intensive than burning fossil fuels. Closely associated with this is the realization that the countries and companies that are at the forefront of developing new energy technologies will have a strong competitive advantage in world markets.

The energy component of Plan B is straightforward. We raise world energy efficiency enough to at least offset all projected growth in energy use from now until 2020. We also turn to wind, solar, geothermal, and other renewable sources to largely replace oil, coal, and natural gas. In effect, Plan B outlines the transition from fossil fuels to renewable sources of energy by 2020. Difficult? Yes. Impossible? No!

Stephen Pacala and Robert Socolow at Princeton University set the stage for Plan B in 2004 when they published an article in *Science* that showed how annual carbon emissions from burning fossil fuels could be held at 7 billion tons instead of rising to 14 billion tons over the next 50 years, as would occur with business as usual. Their goal was to prevent atmospheric CO_2 concentrations, then near 375 ppm, from rising above 500 ppm.

Pacala and Socolow described 15 proven technologies, including efficiency gains and new energy from various renewables, that could each cut carbon emissions 1 billion tons per year by 2054. Any 7 of these options could be combined to prevent an increase in carbon emissions from now through 2054. They further theorized that advancing technology would allow annual carbon emissions to be cut to 2 billion tons by 2104, a level that could likely be absorbed by natural carbon sinks on land and in the oceans.

The Pacala/Socolow exercise was neither a plan nor a projection but a conceptualization, one that has been extraordinarily useful in helping analysts think about the future relationship between energy and climate. Now it is time to select the most promising energy technologies and structure an actual plan to cut carbon emissions. And since climate is changing much faster than anticipated even a few years ago, we believe the world needs to halt the rise in CO_2 levels not at 500 ppm in 2054 but at 400 ppm in 2020. First we look at the enormous potential for raising energy efficiency in the lighting sector.

A REVOLUTION IN LIGHTING TECHNOLOGY

Since the lighting sector is on the edge of a spectacular revolution based on new technologies, perhaps the quickest, most profitable way to reduce electricity use worldwide is simply to change light bulbs.

The first advance in this field came with compact fluorescent lamps (CFLs), which use 75 percent less electricity than old-fashioned incandescents. Replacing inefficient incandescent bulbs that are still widely used today with new

CFLs can reduce the electricity used for lighting by three fourths. Over its lifetime, each standard (13 watt) CFL will reduce electricity bills by roughly $30. And though a CFL may cost twice as much as an incandescent, it lasts 10 times as long. Each one reduces energy use compared with an incandescent by the equivalent of 200 pounds of coal over its lifetime. For perspective, the energy saved by replacing a 100-watt incandescent bulb with an equivalent CFL over its lifetime is sufficient to drive a Toyota Prius hybrid car from New York to San Francisco.

CFL production in China, which accounts for 85 percent of the world total, climbed from 750 million units in 2001 to 2.4 billion units in 2006. Sales in the United States climbed from 21 million CFLs in 2000 to 397 million in 2007. Of the estimated 4.7 billion light sockets in the United States, close to 1 billion now have CFLs.

The world may be moving toward a political tipping point to replace inefficient light bulbs across the board. In February 2007 Australia announced it would phase out the sale of incandescents by 2010, replacing them with CFLs. Canada soon followed with a 2012 phaseout goal. In early 2009, the European Union (EU) approved a phaseout of incandescent bulbs, one that will save the average EU consumer 25–50 euros each year.

Brazil, hit by a nationwide electricity shortage in 2000–02, responded with an ambitious program to replace incandescents with CFLs. As a result, an estimated half of the light sockets there now contain these efficient bulbs. In 2007, China—working with the Global Environment Facility—announced a plan to replace all its incandescents with more-efficient lighting within a decade. And India is planning to phase out incandescent bulbs by 2012.

Retailers are joining the switch too. Wal-Mart, the world's largest retailer, began an ambitious marketing campaign in 2007 to boost its cumulative U.S. sales of compact fluorescents to over 260 million. Currys, Britain's largest electrical retail chain, went further—discontinuing sales of incandescent light bulbs in 2007.

For office buildings, commercial outlets, and factories, where linear (tubular) fluorescents are widely used, the key to cutting electricity use is shifting to the most advanced models, which are even more efficient than CFLs. However, since linear fluorescents are long-lasting, many of those now in use rely on an earlier, less energy-efficient technology.

The second major advance in lighting technology is the light-emitting diode, which uses up to 85 percent less electricity than incandescents. Although LEDs are the ultimate in lighting efficiency, they are still too costly for most uses. They are rapidly taking over several niche markets, however, such as traffic lights, where they now have 52 percent of the U.S. market, and exit signs in buildings, where they hold 88 percent of U.S. sales. New York City has replaced traditional bulbs with LEDs in many of its traffic lights, cutting its annual bill for maintenance and electricity by $6 million. In early 2009, Los Angeles Mayor Antonio Villaraigosa said the city would replace its 140,000 street lights with LEDs, saving taxpayers $48 million over the next seven years. The resulting reduction in carbon emissions would be like taking 7,000 cars off the road.

Universities are also getting involved. In California, the University of California-Davis has a Smart Lighting Initiative. One of its first projects was to replace all the light bulbs in a campus parking garage with LEDs, dramatically reducing electricity use. This success has evolved into LED University, a project to disseminate this technology. Early adopters include the University of California-Santa Barbara, Tianjin Polytechnic University in China, and the University of Arkansas.

LEDs offer another strong economic advantage. While CFLs last 10 times as long as incandescents, LEDs last 50 times as long. Indeed, a typical LED installed at the time of a child's birth will still be working when the youngster graduates from college. The savings in commercial situations from both lower electricity costs and the virtual elimination of replacement maintenance often more than offsets the higher initial cost.

In addition to switching bulbs, energy can be saved just by turning lights off when they are not in use. There are numerous technologies for doing this, including motion sensors that turn lights off in unoccupied offices, living rooms, washrooms, hallways, and stairwells. Sensors and dimmers can also be used to take advantage of daylighting to reduce the intensity of interior lighting when sunlight is bright. In cities, dimmers can be used to reduce streetlight intensity. In fact, these smart lighting technologies can cut the electricity use of LEDs to less than 10 percent of that with incandescents.

In summary, shifting to CFLs in homes, to the most advanced linear fluorescents in office buildings, commercial outlets, and factories, and to LEDs in traffic lights would cut the world share of electricity used for lighting from 19 percent to 7 percent. This would save enough electricity to close 705 of the world's 2,670 coal-fired plants. If the high cost of LEDs drops faster than we have assumed, making widespread use feasible, lighting efficiency gains will come even faster than we have projected.

In a world facing almost daily new evidence of climate change and its consequences, a quick and decisive victory is needed in the battle to cut carbon emissions and stabilize climate. A rapid shift to the most energy-efficient lighting technologies would provide just such a victory—generating momentum for even greater advances in climate stabilization.

ENERGY-EFFICIENT APPLIANCES

Just as CFLs offer great electricity savings over incandescent light bulbs, a similar range of efficiencies is available for many household appliances, such as refrigerators. The U.S. Energy Policy Act of 2005 was designed to exploit some of these potential savings by raising appliance efficiency standards enough to close 29 coal-fired power plants. Other provisions in the act—such as tax incentives that encourage the adoption of energy-efficient technologies, a shift to more combined heat and power generation, and the adoption of real-time pricing of electricity (a measure to discourage optional electricity use during peak demand periods)—would cut electricity demand enough to close an additional 37 coal-fired power plants. Appliance efficiency standards and other measures in the bill would also reduce natural gas consumption substantially. Altogether, these measures are projected to reduce consumer electricity and gas bills in 2020 by more than $20 billion.

Although the U.S. Congress passed legislation raising efficiency for some 30 categories of household and industrial appliances—from refrigerators to industrial-scale electric motors—the U.S. Department of Energy (DOE) has for many years failed to write the standards needed to actually implement the legislation. To remedy this, just days after taking office President Barack Obama ordered the DOE to write regulations to translate law into policy.

With appliances, the big challenge is China. In 1980 its appliance manufacturers produced only 50,000 refrigerators, virtually all for domestic use. In 2008 they produced 48 million refrigerators, 90 million color TVs, and 42 million clothes washers, many of which were for export.

Market penetration of these modern appliances in urban China today is already similar to that in industrial countries. For every 100 urban households there are 138 color TV sets, 97 washing machines, and 88 room air conditioners. Even in rural areas there are 95 color TVs and 46 washing machines for every 100 households. This phenomenal growth in household appliance use in China, along with the extraordinary growth of industry, raised China's electricity use 11-fold from 1980 to 2007. Although China established standards for most appliances by 2005, these are not strictly enforced.

The other major concentration of home appliances is in the European Union, home to 495 million people. Greenpeace notes that even though Europeans on average use half as much electricity as Americans do, they still have a large potential for reducing their usage. A refrigerator in Europe uses scarcely half as much electricity as one in the United States, for example, but the most efficient refrigerators on the market today use only one fourth as much electricity as the

average refrigerator in Europe, suggesting a huge potential for cutting electricity use.

But this is not the end of the efficiency trail, since advancing technology keeps raising the potential. Japan's Top Runner Program is the world's most dynamic system for upgrading appliance efficiency standards. In this system, the most efficient appliances marketed today set the standard for those sold tomorrow. Using this program, between the late 1990s and the end of 2007 Japan raised efficiency standards for individual appliances by anywhere from 15 to 83 percent, depending on the appliance. This is an ongoing process that continually exploits advances in efficiency technologies. A 2008 report indicates that the Top Runner Program for all appliances is running ahead of the ambitious initial expectations—and often by a wide margin.

In an analysis of potential energy savings by 2030 by type of appliance, the Organisation for Economic Co-operation and Development (OECD) put the potential savings from reducing electricity for standby use—the power consumed when an appliance is not being used—at the top of the list. The electricity used by appliances in standby mode worldwide accounts for up to 10 percent of total electricity consumption. In OECD countries, individual household standby power ranged from a low of perhaps 30 watts to a high of over 100 watts in both U.S. and New Zealand households. Since this power is used around the clock, even though the wattage is relatively low, the cumulative use is substantial.

Some governments are capping standby power use by TV sets, computers, microwaves, DVD players, and so on at 1 watt per appliance. South Korea, for example, is mandating a 1-watt limit on standby for many appliances by 2010. Australia is doing the same for nearly all appliances by 2012.

A U.S. study estimates that roughly 5 percent of U.S. residential electricity use is from appliances in standby mode. If this figure dropped to 1 percent, which could be done easily, 17 coal-fired power plants could be closed. If China were to lower its standby losses to 1 percent, it could close a far larger number of power plants.

A more recent efficiency challenge has come with the market invasion of large, flat-screen televisions. The screens now on the market use easily twice as much electricity as a traditional cathode ray tube television. If the flat screen is a large-screen plasma model, it can use four times as much electricity. In the United Kingdom, some Cabinet members are proposing to ban the energy-guzzling flat-screen plasma televisions. California is proposing that all new televisions draw one third less electricity than current sets by 2011 and 49 percent less by 2013.

Consumers often do not buy the most energy-efficient appliances because the initial purchase price is higher, even though this is more than offset by lower appliance lifetime operating costs. If, however, societies adopt a carbon tax reflecting the costs of climate change, the more efficient appliances would be economically much more attractive. Energy use labeling requirements would help consumers choose more wisely.

A worldwide set of appliance efficiency standards keyed to the most efficient models on the market would lead to energy savings in the appliance sector approaching or exceeding the 12 percent of world electricity savings from more-efficient lighting. Thus the combined gains in lighting and appliance efficiencies alone would enable the world to avoid building 1,410 coal-fired power plants—more than the 1,283 new coal-fired power plants that the International Energy Agency (IEA) projects will be built by 2020.

ZERO-CARBON BUILDINGS

The building sector is responsible for a large share of world electricity consumption and raw materials use. In the United States, buildings—commercial and residential—account for 72 percent of electricity use and 38 percent of CO_2 emissions. Worldwide, building construction accounts for 40 percent of materials use.

Because buildings last for 50–100 years or longer, it is often assumed that cutting carbon emissions in the building sector is a long-term process. But that is not the case. An energy retrofit of an older inefficient building can cut energy use

and energy bills by 20–50 percent. The next step, shifting entirely to carbon-free electricity, either generated onsite or purchased, to heat, cool, and light the building completes the job. Presto! A zero-carbon operating building.

The building construction and real estate industries are recognizing what an Australian firm, Davis Langdon, calls "the looming obsolescence of non-green buildings"—one that is driving a wave of reform in both construction and real estate. Further, Davis Langdon says, "going green is future-proofing your asset."

Some countries are taking bold steps. Notable among them is Germany, which as of January 2009 requires that all new buildings either get at least 15 percent of space and water heating from renewable energy or dramatically improve energy efficiency. Government financial support is available for owners of both new and existing buildings for installing renewable energy systems or making efficiency improvements. In reality, once builders or home owners start to plan these installations, they will quickly see that in most cases it makes economic sense to go far beyond the minimal requirements.

There are already signs of progress in the United States. In February 2009, the U.S. Congress passed—and the President signed—the American Recovery and Reinvestment Act, legislation designed to stimulate the U.S. economy. Among other items, it provides for the weatherization of more than a million homes, beginning with an energy audit to identify the measures that would quickly reduce energy use. A second part calls for the weatherization and retrofitting of a large share of the nation's stock of public housing. A third component is the greening of government buildings by making them more energy-efficient and, wherever possible, installing devices such as rooftop solar water and space heaters and rooftop solar electric arrays. The combination of these initiatives is intended to help build a vigorous new industry that would play an active role in raising U.S. energy efficiency and cutting carbon emissions.

In the private sector, the U.S. Green Building Council (USGBC)—well known for its Leadership in Energy and Environmental Design (LEED) certification and rating program—heads the field. This voluntary program, which sets standards well above those of the U.S. government Energy Star building certification program, has four certification levels—certified, silver, gold, and platinum. A LEED-certified building must meet minimal standards in environmental quality, materials use, energy efficiency, and water efficiency. LEED-certified buildings are attractive to buyers because they have lower operating costs, higher lease rates, and typically happier, healthier occupants than traditional buildings do.

The LEED certification standards for construction of new buildings were issued in 2000. Any builder who wants a structure to be rated must request and pay for certification. In 2004 the USGBC also began certifying the interiors of commercial buildings and tenant improvements of existing buildings. And in 2007 it began issuing certification standards for home builders.

Looking at the LEED criteria provides insight into the many ways buildings can become more energy-efficient. The certification process for new buildings begins with site selection, and then moves on to energy efficiency, water efficiency, materials use, and indoor environmental quality. In site selection, points are awarded for proximity to public transport, such as subway, light rail, or bus lines. Beyond this, a higher rating depends on provision of bicycle racks and shower facilities for employees. New buildings must also maximize the exposure to daylight, with minimum daylight illumination for 75 percent of the occupied space.

With energy, exceeding the high level of efficiency required for basic certification earns additional points. Further points are awarded for the use of renewable energy, including rooftop solar cells to generate electricity, rooftop solar water and space heaters, and the purchase of green power.

Thus far LEED has certified 1,600 new buildings in the United States, with some 11,600 planned or under construction that have applied for certification. The commercial building space that has been certified or registered for certification approval totals 5 billion square feet of floor

space, or some 115,000 acres (the equivalent of 115,000 football fields).

The Chesapeake Bay Foundation's office building for its 100 staff members near Annapolis, Maryland, was the first to earn a LEED platinum rating. Among its features are a ground-source heat pump for heating and cooling, a rooftop solar water heater, and sleekly designed composting toilets that produce a rich humus used to fertilize the landscape surrounding the building.

Toyota's North American headquarters in Torrance, California, which houses 2,000 employees, has a LEED gold rating and is distinguished by a large solar-electric generating facility that provides much of its electricity. Waterless urinals and rainwater recycling enable it to operate with 94 percent less water than a conventionally designed building of the same size. Less water use also means less energy use.

The 54-story Bank of America tower in New York is the first large skyscraper expected to earn a platinum rating. It has its own co-generation power plant and collects rainwater, reuses waste water, and used recycled materials in construction.

A 60-story office building with a gold rating being built in Chicago will use river water to cool the building in summer, and the rooftop will be covered with plants to reduce runoff and heat loss. Energy-conserving measures will save the owner $800,000 a year in energy bills. The principal tenant, Kirkland and Ellis LLP, a Chicago-based law firm, insisted that the building be gold-certified and that this be incorporated into the lease.

The state of California commissioned Capital E, a green building consulting firm, to analyze the economics of 33 LEED-certified buildings in the state. The study concluded that certification raised construction costs by $4 per square foot but that because operating costs as well as employee absenteeism and turnover were lower and productivity was higher than in other buildings, the standard- and silver-certified buildings earned a profit over the first 20 years of $49 per square foot, and the gold- and platinum-certified buildings earned $67 per square foot.

In 2002 a global version of the USGBC, the World Green Building Council, was formed. As of spring 2009 it included Green Building Councils in 14 countries, including Brazil, India, and the United Arab Emirates. Eight other countries—ranging from Spain to Vietnam—are working to meet the prerequisites for membership. Among the current members, India ranks second in certification after the United States, with 292 million square feet of LEED-certified floor space, followed by China (287 million) and Canada (257 million).

Beyond greening new buildings, there are numerous efforts to make older structures more efficient. In 2007, the Clinton Foundation announced an Energy Efficiency Building Retrofit Program, a project of the Clinton Climate Initiative. In cooperation with C40, a large-cities climate leadership group, this program brings together five of the world's largest banks and four leading energy service companies to work with an initial group of 16 cities to retrofit buildings, reducing their energy use by 20–50 percent. Among the cities are some of the world's largest: Bangkok, Berlin, Karachi, London, Mexico City, Mumbai, New York, Rome, and Tokyo. Each of the banks involved—ABN AMRO, Citi, Deutsche Bank, JP Morgan Chase, and UBS—is committed to investing up to $1 billion in this effort, enough to easily double the current worldwide level of energy saving retrofits.

The energy service companies—Honeywell, Johnson Controls, Siemens, and Trane—committed not only to do the actual retrofitting but also to provide "performance guarantees," thus ensuring that all the retrofits will be profitable. At the launch of this program, former President Bill Clinton pointed out that banks and energy service companies would make money, building owners would save money, and carbon emissions would fall. As of February 2009, the Clinton Climate Initiative had been involved with 250 retrofit projects and over 500 million square feet of floor space.

In April 2009, the owners of New York's Empire State Building announced plans to retrofit the 2.6 million square feet of office space in the nearly 80-year-old 102-story building, thereby

reducing its energy use by nearly 40 percent. The resulting energy savings of $4.4 million a year is expected to recover the retrofitting costs in three years.

Beyond these voluntary measures, the government-designed building codes, which set minimal standards for building energy efficiency, are highly effective. In the United States this has been dramatically demonstrated in differences between California and the country as a whole in housing energy efficiency. Between 1975 and 2002, residential energy use per person dropped 16 percent in the country as a whole. But in California, which has stringent building codes, it dropped by 40 percent. The bottom line is that there is an enormous potential for reducing energy use in buildings in the United States and, indeed, the world.

One firm believer in that potential is Edward Mazria, a climate-conscious architect from New Mexico. He has launched the 2030 Challenge. Its principal goal is for U.S. architects to be designing buildings in 2030 that use no fossil fuels. Mazria observes that the buildings sector is the leading source of carbon emissions, easily eclipsing transportation. Therefore, he says, "it's the architects who hold the key to turning down the global thermostat." To reach his goal, Mazria has organized a coalition of several organizations, including the American Institute of Architects, the USGBC, and the U.S. Conference of Mayors.

Mazria also recognizes the need for faculty retraining in the country's 124 architectural schools to "transform architecture from its mindless and passive reliance on fossil fuels to an architecture intimately linked to the natural world in which we live."

Today's architectural concepts and construction technologies enable architects to easily design new buildings with half the energy requirements of existing ones. Among the design technologies they can use are natural daylighting, rooftop solar-electric cells, rooftop solar water and space heaters, ultra insulation, natural ventilation, ground source heat pumps, glazed windows, waterless urinals, more-efficient lighting technologies, and motion sensors for lighting. Designing and constructing energy-efficient buildings, combined with a massive harnessing of renewable energy, makes it not only possible but also profitable for buildings to operate without fossil fuels.

ELECTRIFYING THE TRANSPORT SYSTEM

Among the keys to cutting carbon emissions are redesigning urban transport and the overall electrification of transportation. The last century witnessed the evolution of an oil-powered transport system: gasoline for cars and diesel fuel for trucks and trains. Now that is changing. With both cars and rail systems, oil will be replaced by electricity. And the power will come increasingly from wind farms and from solar and geothermal power plants.

With peak oil on our doorstep, the world desperately needs a new automotive energy economy. Fortunately, the foundation for this has been laid with two new technologies: gas-electric hybrid plug-in cars and all-electric cars.

The Toyota Prius—the world's top-selling hybrid car—gets an impressive 50 miles per gallon (mpg) in combined city/highway driving, nearly double that of the average new U.S. passenger vehicle. The United States could easily cut its gasoline use in half simply by converting the entire American automobile fleet to highly efficient hybrid cars. But this is only the beginning.

Now that hybrid cars are well established, it is a relatively small additional step to manufacture plug-in hybrids that run largely on electricity. By shifting to lithium ion batteries to boost electricity storage capacity and by adding an extension cord so the battery can be recharged from the grid, drivers can do their commuting, grocery shopping, and other short-distance travel almost entirely with electricity, using gasoline only for the occasional long trip. Even more exciting, recharging batteries with off-peak wind-generated electricity costs the equivalent of less than $1 per gallon of gasoline.

As of mid-2009, nearly all major car makers have announced plans to bring either plug-in hybrids or all-electric cars to market. The world's first commercially available plug-in hybrid car

reached the market in December 2008 in China. While the world was focusing on the race between Toyota and GM, China's BYD (Build Your Dreams) had quietly forged ahead, bringing its plug-in hybrid car to market. Already in mass production and selling for a highly competitive $22,000, it is scheduled to appear in U.S. and European markets in 2010.

Meanwhile, Toyota apparently has gotten the jump on GM by announcing it would start to market a limited number of plug-in hybrids for selected use by the end of 2009. The Chevrolet Volt, GM's entry, is expected to average 150 mpg, largely because of a stronger battery and greater all-electric range. It is this prospect of triple-digit gasoline mileage that is selling customers on plug-in hybrids.

Nissan has been emphasizing the development of an all-electric car, which it plans to market in 2010. Chrysler plans to produce an electric version of several of its models, effectively offering customers a choice between gasoline and electrically powered vehicles. Think, an entrepreneurial Norwegian firm, already producing an all-electric car in Norway, is planning an assembly plant in the United States in 2010 to produce up to 60,000 electric cars per year.

Shifting to plug-in electric hybrids and all-electric cars does not require a costly new infrastructure, since the network of gasoline service stations and the electricity grid are already in place. A 2006 study by the U.S. Pacific Northwest National Laboratory estimated that over 80 percent of the electricity needs of a national fleet of all plug-in cars could be satisfied with the existing electrical infrastructure since the recharging would take place largely at night, when there is an excess of generating capacity. What will be needed is the installation of electrical outlets in parking garages, parking lots, and street-side parking meters, along with a credit card access device to identify the user for billing purposes.

Silicon Valley entrepreneur Shai Agassi is working with Nissan and the governments of Israel, Denmark, Australia, and Canada's Ontario Province, as well as the San Francisco Bay area of California and Hawaii in the United States, to set up networks of electric-car service stations. These stations would replace depleted batteries with freshly charged ones, thus eliminating the need for time-consuming recharges. Whether the typical daily driving distance will warrant investment in battery replacement on this scale remains to be seen.

While the future of transportation in cities lies with a mix of light rail, buses, bicycles, some cars, and walking, the future of intercity travel belongs to high-speed trains. Japan, with its bullet trains, pioneered this mode of travel. Operating at speeds up to 190 miles per hour, Japan's bullet trains carry almost a million passengers a day. On some of the heavily used intercity high-speed lines, trains depart every three minutes.

Beginning in 1964 with the 322-mile line from Tokyo to Osaka, Japan's high-speed rail network now stretches for 1,360 miles, linking nearly all its major cities. One of the most heavily traveled links is the original line, where the bullet trains carry 413,000 passengers a day. The transit time of two-and-a-half hours between Tokyo and Osaka compares with a driving time of eight hours. High-speed trains save time as well as energy.

Although Japan's bullet trains have carried billions of passengers in great comfort over 40 years at high speeds, there has not been a single casualty. Late arrivals average 6 seconds. If we were selecting seven wonders of the modern world, Japan's high-speed rail system surely would be among them.

Although the first European high-speed line, from Paris to Lyon, did not begin operation until 1981, Europe has made enormous strides since then. As of 2009 there were 3,100 miles (5,000 kilometers) of high-speed rail operating in Europe. The goal is to triple this track length by 2020 and eventually to integrate the eastern countries, including Poland, the Czech Republic, and Hungary, into a continental network.

While France and Germany were the early European leaders in intercity rail, Spain is fast building a high-speed intercity rail network as well. Within a year of opening the Barcelona-to-Madrid connection, domestic airlines lost roughly a fifth of their passengers to these high-speed intercity trains. Spain plans to link with

high-speed systems in France to become firmly integrated into the European network.

Existing international links, such as the one between Paris and Brussels, are being joined by connections between Paris and Stuttgart, Frankfurt and Paris, and London and Paris (the latter via the Channel Tunnel). On the newer lines, trains are operating at up to 200 miles per hour. As *The Economist* notes, "Europe is in the grip of a high speed rail revolution."

High-speed links between cities dramatically raise rail travel. For example, when the Paris-to-Brussels link opened—the 194 miles is covered by train in just 85 minutes—the share of those traveling between the two cities by train rose from 24 percent to 50 percent. The car share dropped from 61 to 43 percent, and plane travel virtually disappeared.

Carbon dioxide emissions per passenger mile on electric high-speed trains are roughly one third those of cars and one fourth those of planes. In the Plan B economy, carbon emissions from trains will essentially be zero, since they will be powered almost entirely by renewable electricity. In addition to being comfortable and convenient, these rail links reduce air pollution, congestion, noise, and accidents. They also free travelers from the frustrations of traffic congestion and long airport security check lines.

There is a huge gap in high-speed rail between Japan and Europe on the one hand and the rest of the world on the other. But China is beginning to develop high-speed trains linking some of its major cities. A high-speed link between Beijing and Shanghai scheduled for completion by 2013 will slice train travel time in half, from 10 to 5 hours. China now has 3,890 miles of track that can handle train speeds of up to 125 miles per hour. The plan is to triple the length of high-speed track by 2020.

The United States has a "high-speed" Acela Express that links Washington, New York, and Boston, but unfortunately neither its rail bed and speed nor its reliability come close to the trains in Japan and Europe. The good news is that the U.S. economic stimulus plan signed into law in February 2009 contained some $8 billion to help launch a new era of high-speed rail construction in the United States.

In the United States, the need to cut carbon emissions and prepare for shrinking oil supplies calls for this shift in investment from roads and highways to railways. In 1956 U.S. President Dwight Eisenhower launched the interstate highway system, justifying it on national security grounds. Today the threat of climate change and the insecurity of oil supplies argue for the construction of a high-speed electrified rail system, for both passenger and freight traffic. The additional electricity needed could easily be supplied from renewable sources, mainly wind energy.

The passenger rail system would be modeled after those of Japan and Europe. A high-speed transcontinental line that averaged 170 miles per hour would mean traveling coast-to-coast in 15 hours, even with stops in major cities along the way. There is a parallel need to develop an electrified national rail freight network that would greatly reduce the need for long-haul trucks.

Voters in California approved a bond referendum in November 2008 of nearly $10 billion to build a high-speed rail system to link northern and southern California. This would reduce the use of cars and eliminate many of the fuel-guzzling short-distance flights linking California's major cities.

Any meaningful global effort to cut transport carbon emissions begins with the United States, which consumes more gasoline than the next 20 countries combined (including Japan, China, Russia, Germany, and Brazil). The United States—with 249 million passenger vehicles out of the global 912 million— not only has the largest fleet but it is near the top in miles driven per car and near the bottom in fuel efficiency.

The first step to reduce this massive U.S. consumption of gasoline is to raise fuel efficiency standards. The 40-percent increase in these standards by 2016 announced by the Obama administration in May 2009 will greatly reduce U.S. gasoline use and bring the country closer to the fuel economy levels prevailing in Europe and China. A crash program to shift the U.S. fleet to plug-in hybrids and all-electric cars would make

an even greater contribution. And shifting public funds from highway construction to public transit would reduce the number of cars needed, bringing us close to our goal of cutting carbon emissions 80 percent by 2020.

A NEW MATERIALS ECONOMY

The production, processing, and disposal of materials in our modern throwaway economy wastes not only materials but energy as well. In nature, one-way linear flows do not survive long. Nor, by extension, can they survive long in the expanding global economy. The throwaway economy that has evolved over the last half-century is an aberration, now itself headed for the junk heap of history.

The potential for sharply reducing materials use was first identified in Germany, initially by Friedrich Schmidt-Bleek in the early 1990s and then by Ernst von Weizsäcker, an environmental leader in the German Bundestag. They argued that modern industrial economies could function very effectively using only one fourth the virgin raw materials prevailing at the time. A few years later, Schmidt-Bleek, who founded the Factor Ten Institute in France, showed that raising resource productivity even more—by a factor of 10—was well within the reach of existing technology and management, given the right policy incentives.

In their book *Cradle to Cradle: Remaking the Way We Make Things*, American architect William McDonough and German chemist Michael Braungart conclude that waste and pollution are to be avoided entirely. "Pollution," says McDonough, "is a symbol of design failure."

Beyond reducing materials use, the energy savings from recycling are huge. Steel made from recycled scrap takes only 26 percent as much energy as that from iron ore. For aluminum, the figure is just 4 percent. Recycled plastic uses only 20 percent as much energy. Recycled paper uses 64 percent as much—and with far fewer chemicals during processing. If the world recycling rates of these basic materials were raised to those already attained in the most efficient economies, carbon emissions would drop precipitously.

Industry, including the production of plastics, fertilizers, steel, cement, and paper, accounts for more than 30 percent of world energy consumption. The petrochemical industry, which produces such things as plastics, fertilizer, and detergents, is the biggest consumer of energy in the manufacturing sector, accounting for about a third of worldwide industrial energy use. Since a large part of industry fossil fuel use is for feedstock to manufacture plastics and other materials, increased recycling can reduce feedstock needs. Worldwide, increasing recycling rates and moving to the most efficient manufacturing systems in use today could easily reduce energy use in the petrochemical industry by 32 percent.

The global steel industry, producing over 1.3 billion tons in 2008, accounts for 19 percent of industrial energy use. Efficiency measures, such as adopting the most efficient blast furnace systems in use today and the complete recovery of used steel, could reduce energy use in the steel industry by 23 percent.

Reducing materials use begins with recycling steel, the use of which dwarfs that of all other metals combined. Steel use is dominated by three industries—automobile, household appliances, and construction. In the United States, virtually all cars are recycled. They are simply too valuable to be left to rust in out-of-the-way junkyards. The U.S. recycling rate for household appliances is estimated at 90 percent. For steel cans it is 63 percent, and for construction steel the figures are 98 percent for steel beams and girders but only 65 percent for reinforcement steel. Still, the steel discarded each year in various forms is enough to meet the needs of the U.S. automobile industry.

Steel recycling started climbing more than a generation ago with the advent of the electric arc furnace, a technology that produces steel from scrap using only one fourth the energy required to produce it from virgin ore. Electric arc furnaces using scrap now account for half or more of steel production in more than 20 countries. A few countries, including Venezuela and Saudi Arabia, use electric arc furnaces exclusively. If three fourths of steel production were to switch to electric arc furnaces using scrap, energy use in

the steel industry could be cut by almost 40 percent.

The cement industry, turning out 2.9 billion tons in 2008, is another major energy consumer. China, accounting for half of world production, manufactures more cement than the next 20 countries combined, yet it does so with extraordinary inefficiency. If China used the same kiln technologies as Japan, it could reduce its cement production energy use by 45 percent. Worldwide, if all cement producers used the most efficient dry kiln process, energy use in the industry could drop 42 percent.

Restructuring the transportation system also has a huge potential for reducing materials use as light rail and buses replace cars. For example, improving urban transit means that one 12-ton bus can easily replace 60 cars weighing 1.5 tons each, or a total of 90 tons, reducing material use 87 percent. And every time someone replaces a car with a bike, material use is reduced 99 percent.

The big challenge for cities in saving energy is to recycle as many components of the urban waste flow as possible. Virtually all paper products can now be recycled, including cereal boxes, junk mail, and paper bags in addition to newspapers and magazines. So too can metals, glass, and most plastics. Kitchen and yard waste can be composted to recycle plant nutrients.

Advanced industrial economies with stable populations, such as those in Europe and Japan, can rely primarily on the stock of materials already in the economy rather than using virgin raw materials. Metals such as steel and aluminum can be used and reused indefinitely.

In the United States, the latest *State of Garbage in America* report shows that 29 percent of garbage is recycled, 7 percent is burned, and 64 percent goes to landfills. Recycling rates among U.S. cities vary from less than 30 percent in some cities to more than 70 percent in San Francisco, the highest in the country. When San Francisco hit 70 percent in 2008, Mayor Gavin Newsom immediately announced a plan to reach 75 percent. Among the largest U.S. cities, recycling rates vary from 34 percent in New York to 55 percent in Chicago and 60 percent in Los Angeles. At the state level, Florida has boldly set a goal of recycling 75 percent of waste by 2020.

One of the most effective ways to encourage recycling is to adopt a landfill tax. For example, when the state of New Hampshire adopted a "pay-as-you-throw" program that encourages municipalities to charge residents for each bag of garbage, it dramatically reduced the flow of materials to landfills. In the small town of Lyme, with nearly 2,000 people, adoption of a landfill tax raised the share of garbage recycled from 13 to 52 percent in one year.

The recycled material in Lyme, which jumped from 89 tons in 2005 to 334 tons in 2006, included corrugated cardboard, which sold for $90 a ton, mixed paper at $45 a ton, and aluminum at $1,500 a ton. This program simultaneously reduced the town's landfill fees while generating a cash flow from the sale of recycled material.

In addition to measures that encourage recycling, there are those that encourage or mandate the reuse of products such as beverage containers. Finland, for example, has banned the use of one-way soft drink containers. Canada's east coast province, Prince Edward Island, has adopted a similar ban on all nonrefillable beverage containers. The result in both cases is a sharply reduced flow of garbage to landfills. A refillable glass bottle used over and over requires about 10 percent as much energy per use as an aluminum can that is recycled. Cleaning, sterilizing, and relabeling a used bottle requires little energy compared with recycling cans made from aluminum, which has a melting point of 1,220 degrees Fahrenheit. Banning nonrefillables is a quintuple win option—cutting material use, carbon emissions, air pollution, water pollution, and landfill costs simultaneously. There are also substantial transport fuel savings, since the refillable containers are simply back-hauled by delivery trucks to the original bottling plants or breweries for refilling.

San José, California, already diverting 62 percent of its municipal waste from landfills for reuse and recycling, is now focusing on the large flow of trash from construction and demolition sites. This material is trucked to one of two dozen specialist recycling firms in the city. For example, at Premier

Recycle up to 300 tons of building debris are delivered each day. This is skillfully separated into recyclable piles of concrete, scrap metal, wood, and plastics. Some materials the company sells, some it gives away, and some it just pays someone to take.

Before the program began, only about 100,000 tons per year of San José's mixed construction and demolition materials were reused or recycled. Now it is nearly 500,000 tons. The scrap metal that is salvaged goes to recycling plants, wood can be converted into gardening mulch or into wood chips for fueling power plants, and concrete can be recycled to build road banks. By deconstructing a building instead of simply demolishing it, most of the material in it can be reused or recycled, thus dramatically reducing energy use and carbon emissions. San José is becoming a model for cities everywhere.

Germany and, more recently, Japan are requiring that products such as automobiles, household appliances, and office equipment be designed for easy disassembly and recycling. In May 1998, the Japanese Diet enacted a tough appliance recycling law, one that prohibits discarding household appliances, such as washing machines, TV sets, or air conditioners. With consumers bearing the cost of disassembling appliances in the form of a disposal fee to recycling firms, which can come to $60 for a refrigerator or $35 for a washing machine, there is strong pressure to design appliances so they can be more easily and cheaply disassembled.

With computers becoming obsolete every few years as technology advances, the need to quickly disassemble and recycle them is another paramount challenge in building an eco-economy. In Europe, information technology (IT) firms are exploring the reuse of computer components. Because European law requires manufacturers to pay for the collection, disassembly, and recycling of toxic materials in IT equipment, they have begun to focus on how to disassemble everything from computers to cell phones. Finland-based Nokia, for example, has designed a cell phone that will virtually disassemble itself.

On the clothing front, Patagonia, an outdoor gear retailer, has launched a garment recycling program beginning with its polyester fiber garments. Working with Teijin, a Japanese firm, Patagonia is taking back and recycling not only the polyester garments it sells but also those sold by its competitors. Patagonia estimates that making a garment from recycled polyester, which is indistinguishable from the initial polyester made from petroleum, uses less than one fourth as much energy. With this success behind it, Patagonia has broadened the program to recycle its cotton tee shirts as well as nylon and wool clothing.

Remanufacturing is even more efficient. Within the heavy industry sector, Caterpillar has emerged as a leader. At a plant in Corinth, Mississippi, the company recycles some 17 truckloads of diesel engines a day. These engines, retrieved from Caterpillar's clients, are disassembled by hand by workers who do not throw away a single component, not even a bolt or screw. Once the engine is disassembled, it is reassembled with all worn parts repaired or replaced. The resulting engine is as good as new. In 2006, Caterpillar's remanufacturing division was racking up $1 billion a year in sales and growing at 15 percent annually, contributing impressively to the company's bottom line.

Another emerging industry is airliner recycling. Daniel Michaels writes in the *Wall Street Journal* that Boeing and Airbus, which have been building jetliners in competition for nearly 40 years, are now vying to see who can dismantle planes most efficiently. The first step is to strip the plane of its marketable components, such as engines, landing gear, galley ovens, and hundreds of other items. For a jumbo jet, these key components can collectively sell for up to $4 million. Then comes the final dismantling and recycling of aluminum, copper, plastic, and other materials. The next time around the aluminum may show up in cars, bicycles, or another jetliner.

The goal is to recycle 90 percent of the plane, and perhaps one day 95 percent or more. With more than 3,000 airliners already put out to pasture and many more to come, this retired fleet has become the equivalent of an aluminum mine.

Another increasingly attractive option for cutting carbon emissions is to discourage energy-intensive but nonessential industries. The gold jewelry, bottled water, and plastic bag industries are prime examples. The annual world production of 2,380 tons of gold, the bulk of it used for jewelry, requires the processing of 500 million tons of ore. For comparison, while 1 ton of steel requires the processing of 2 tons of ore, 1 ton of gold involves processing an almost incomprehensible 200,000 tons of ore. Processing ore for gold consumes a vast amount of energy—and emits as much CO_2 as 5.5 million cars.

In a world trying to stabilize climate, it is very difficult to justify bottling water (often tap water to begin with), hauling it over long distances, and then selling it for 1,000 times the price of tap water. Although clever marketing, designed to undermine public confidence in the safety and quality of municipal water supplies, has convinced many consumers that bottled water is safer and healthier than water from faucets, a detailed study by the World Wide Fund for Nature could not find any support for this claim. It notes that in the United States and Europe there are more standards regulating the quality of tap water than bottled water. For people in developing countries where water is unsafe, it is far cheaper to boil or filter water than to buy it in bottles.

Manufacturing the nearly 28 billion plastic bottles used each year to package water in the United States alone requires the equivalent of 17 million barrels of oil. And whereas tap water is delivered through a highly energy-efficient infrastructure, bottled water is hauled by trucks, sometimes over hundreds of miles. Including the energy for hauling water from bottling plants to sales outlets and the energy needed for refrigeration, the U.S. bottled water industry consumes roughly 50 million barrels of oil per year, enough oil to fuel 3 million cars for one year.

The good news is that people are beginning to see how wasteful and climate-disruptive this industry is. Mayors of U.S. cities are refusing to spend taxpayer dollars to buy bottled water for their employees at exorbitant prices when high-quality tap water is readily available. Mayor Rocky Anderson of Salt Lake City noted the "total absurdity and irresponsibility, both economic and environmental, of purchasing and using bottled water when we have perfectly good and safe sources of tap water."

San Francisco Mayor Newsom has banned the use of city funds to purchase bottled water. Other cities following a similar strategy include Los Angeles, Salt Lake City, and St. Louis. New York City has launched a $5-million ad campaign to promote its tap water and thus to rid the city of bottled water and the fleets of delivery trucks that tie up traffic. In response to initiatives such as these, U.S. sales of bottled water began to decline in 2008.

Like plastic water bottles, throwaway plastic shopping bags are also made from fossil fuels, can take centuries to decompose, and are almost always unnecessary. In addition to local initiatives, several national governments are moving to ban or severely restrict the use of plastic shopping bags, including China, Ireland, Eritrea, Tanzania, and the United Kingdom.

In summary, there is a vast worldwide potential for cutting carbon emissions by reducing materials use. This begins with the major metals—steel, aluminum, and copper—where recycling requires only a fraction of the energy needed to produce these metals from virgin ore. It continues with the design of cars, household appliances, and electronic products so they are easily disassembled into their component parts for reuse or recycling. And it includes avoiding unnecessary products.

SMARTER GRIDS, APPLIANCES, AND CONSUMERS

More and more utilities are beginning to realize that building large power plants just to handle peak daily and seasonal demand is a very costly way of managing an electricity system. Existing electricity grids are typically a patchwork of local grids that are simultaneously inefficient, wasteful, and dysfunctional in that they often are unable, for example, to move electricity surpluses to areas of shortages. The U.S. electricity grid today resembles the roads and highways of the mid-twentieth century before the interstate highway system was built. What is needed today is the

electricity equivalent of the interstate highway system.

The inability to move low-cost electricity to consumers because of congestion on transmission lines brings with it costs similar to those associated with traffic congestion. The lack of transmission capacity in the eastern United States is estimated to cost consumers $16 billion a year in this region alone.

In the United States, a strong national grid would permit power to be moved continuously from surplus to deficit regions, thus reducing the total generating capacity needed. Most important, the new grid would link regions rich in wind, solar, and geothermal energy with consumption centers. A national grid, drawing on a full range of renewable energy sources, would itself be a stabilizing factor.

Establishing strong national grids that can move electricity as needed and that link new energy sources with consumers is only half the battle, however. The grids and appliances need to become "smarter" as well. In the simplest terms, a smart grid is one that takes advantage of advances in information technology, integrating this technology into the electrical generating, delivery, and user system, enabling utilities to communicate directly with customers and, if the latter agree, with their household appliances.

Smart grid technologies can reduce power disruption and fluctuation that cost the U.S. economy close to $100 billion a year, according to the Electric Power Research Institute. In an excellent 2009 Center for American Progress study, *Wired for Progress 2.0: Building a National Clean-Energy Smart Grid*, Bracken Hendricks notes the vast potential for raising grid efficiency with several information technologies: "A case in point would be encouraging the widespread use of synchrophasors to monitor voltage and current in real time over the grid network. It has been estimated that better use of this sort of real-time information across the entire electrical grid could allow at least a 20 percent improvement in energy efficiency in the United States." This and many other examples give us a sense of the potential for increasing grid efficiency.

A smart grid not only moves electricity more efficiently in geographic terms; it also enables electricity use to be shifted over time—for example, from periods of peak demand to those of off-peak demand. Achieving this goal means working with consumers who have "smart meters" to see exactly how much electricity is being used at any particular time. This facilitates two-way communication between utility and consumer so they can cooperate in reducing peak demand in a way that is advantageous to both. And it allows the use of two-way metering so that customers who have a rooftop solar electric panel or their own windmill can sell surplus electricity back to the utility.

Smart meters coupled with smart appliances that can receive signals from the grid allow electricity use to be shifted away from peak demand. Higher electricity prices during high demand periods also prod consumers to change their behavior, thus improving market efficiency. For example, a dishwasher can be programmed to run not at 8 p.m. but at 3 a.m., when electricity demand is much lower, or air conditioners can be turned off for a brief period to lighten the demand load.

Another approach being pioneered in Europe achieves the same goal but uses a different technology. In any grid, there is a narrow range of fluctuation in the power being carried. An Italian research team is testing refrigerators that can monitor the grid flow and, when demand rises or supply drops, simply turn themselves off for as long as it is safe to do so. *New Scientist* reports that if this technology were used in the 30 million refrigerators in the United Kingdom, it would reduce national peak demand by 2,000 megawatts of generating capacity, allowing the country to close four coal-fired power plants.

A similar approach could be used for air conditioning systems in both residential and commercial buildings. Karl Lewis, COO of GridPoint, a U.S. company that designs smart grids, says "we can turn off a compressor in somebody's air conditioning system for 15 minutes and the temperature really won't change in the house." The bottom line with a smart grid is that a modest investment in information technology can reduce peak power, yielding both savings in electricity

and an accompanying reduction in carbon emissions.

Some utilities are pioneers in using time-based pricing of electricity, when electricity used during off-peak hours is priced much lower than that used during peak hours. Similarly, in regions with high summer temperatures, there is often a costly seasonal peak demand. Baltimore Gas and Electric (BGE), for example, conducted a pilot program in 2008 in which participating customers who permitted the utility to turn off their air conditioners for selected intervals during the hottest days were credited generously for the electricity they saved. The going rate in the region is roughly 14¢ per kilowatt-hour. But for a kilowatt-hour saved during peak hours on peak days, customers were paid up to $1.75—more than 12 times as much. Thus if they saved 4 kilowatt-hours of electricity in one afternoon, they got a $7 credit on their electricity bill. Customers reduced their peak electricity consumption by as much as one third, encouraging BGE to design a similar program with even more "smart" technology for summer 2009.

Within the United States the shift to smart meters is moving fast, with some 28 utilities planning to deploy smart meters in the years ahead. Among the leaders are California's two major utilities, Pacific Gas and Electric and Southern California Edison, which are planning on full deployment to their 5.1 million and 5.3 million customers by 2012. Both will offer variable rates to reduce peak electricity use. Among the many other utilities aiming for full deployment are American Electric Power in the Midwest (5 million customers) and Florida Power and Light (4.4 million customers).

Europe, too, is installing smart meters, with Finland setting the pace. A Swedish research firm, Berg Insight, projects that Europe will have 80 million smart meters installed by 2013.

Unfortunately, the term "smart meters" describes a wide variety of meters, ranging from those that simply provide consumers with real-time data on electricity use to those that facilitate two-way communication between the utility and customer or even between the utility and individual household appliances. The bottom line: the smarter the meter, the greater the savings.

Taking advantage of information technology to increase the efficiency of the grid, the delivery system, and the use of electricity at the same time is itself a smart move. Simply put, a smart grid combined with smart meters enables both electrical utilities and consumers to be much more efficient.

THE ENERGY SAVINGS POTENTIAL

The goal for this chapter was to identify energy-saving measures that will offset the nearly 30 percent growth in global energy demand projected by the IEA between 2006 and 2020. My colleagues and I are confident that the measures proposed will more than offset the projected growth in energy use.

Shifting to more energy-efficient lighting alone lowers world electricity use by 12 percent. With appliances, the key to raising energy efficiency is to establish international efficiency standards that reflect the most efficient models on the market today, regularly raising this level as technologies advance. This would in effect be the global version of Japan's Top Runner program to raise appliance efficiency.

Given the potential for raising appliance efficiency, the energy saved by 2020 should at least match the savings in the lighting sector. Combining more-efficient lights and appliances with a smart grid that uses time-of-day pricing, peak demand sensors, and the many other technologies described in this chapter shows a huge potential for reducing both overall electricity use and peak demand.

It is easy to underestimate the potential for reducing electricity use. Within the United States, the Rocky Mountain Institute calculates that if the 40 least efficient states were to achieve the electrical efficiency of the 10 most efficient ones, national electricity use would be cut by one third. This would allow the equivalent of 62 percent of all U.S. coal-fired power plants to be closed down. But even the most efficient states have a substantial potential for further reducing

electricity use and, indeed, are planning to keep cutting carbon emissions and saving money.

In terms of transportation, the short-term keys to reducing oil use and carbon emissions involve shifting to highly fuel-efficient cars (including electric vehicles), diversifying urban transport systems, and building intercity rapid rail systems modeled on those in Japan and Europe. This shift from car-dominated transport systems to diversified systems is evident in the actions of hundreds of mayors worldwide who struggle daily with traffic congestion and air pollution. They are devising ingenious ways of limiting not only the use of cars but also the very need for them. As the urban car presence diminishes, the nature of the city itself will change.

Within the industrial sector, there is a hefty potential for reducing energy use. In the petrochemical industry, moving to the most efficient production technologies now available and recycling more plastic can cut energy use by 32 percent. Gains in manufacturing efficiency in steel can cut energy use by 23 percent. Even larger gains are within reach for cement, where simply shifting to the most efficient dry kiln technologies can reduce energy use by 42 percent.

With buildings—even older buildings, where retrofitting can reduce energy use by 20–50 percent—there is a profitable potential for saving energy. As noted earlier, such a reduction in energy use, combined with the use of renewable electricity to heat, cool, and light the building, means that it will be easier to create carbon-neutral buildings than we may have thought.

One simple way to achieve all these gains is to adopt a carbon tax that would help reflect the full cost of burning fossil fuels. We recommend increasing this carbon tax by $20 per ton each year over the next 10 years, for a total of $200 ($55 per ton of CO_2), offsetting it with a reduction in income taxes. High though this may seem, it does not come close to covering all the indirect costs of burning fossil fuels. It does, however, encourage investment in both efficiency and carbon-free sources of energy.

In seeking to raise energy efficiency as described in this chapter, there have been some exciting surprises in the vast potential for doing so. We now turn to developing the earth's renewable sources of energy, where there are equally exciting possibilities.